



BANK OF ENGLAND

Staff Working Paper No. 951

Unlocking new methods to estimate country-specific trade costs and trade elasticities

Rebecca Freeman, Mario Larch, Angelos Theodorakopoulos
and Yoto Yotov

November 2021

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Committee.



BANK OF ENGLAND

Staff Working Paper No. 951

Unlocking new methods to estimate country-specific trade costs and trade elasticities

Rebecca Freeman,⁽¹⁾ Mario Larch,⁽²⁾ Angelos Theodorakopoulos⁽³⁾ and Yoto Yotov⁽⁴⁾

Abstract

We propose new methods to identify the full impact of country-specific characteristics on bilateral trade flows within the framework of 'the new quantitative trade model'. We complement theory with a simple two-stage estimating procedure, and offer a proof of concept by quantifying the impact of country-specific research and development expenditure on trade. Results suggest a positive relationship overall, but a larger impact on international (versus domestic) trade. Further, our methodology allows us to recover trade elasticity estimates without the need for price/tariff data. Bringing this to the sectoral level, we obtain estimates of the trade elasticity for manufacturing, services, and tradable versus non-tradable sectors.

Key words: Structural gravity, country-specific trade costs, trade elasticity, elasticity of substitution, R&D, trade.

JEL classification: F10, F14, F16.

(1) Bank of England and CEP/LSE. Email: rebecca.freeman@bankofengland.co.uk

(2) Department of Law and Economics, University of Bayreuth, Germany; CEPII; CESifo; ifo; and GEP.
Email: mario.larch@uni-bayreuth.de

(3) University of Oxford, Oxford Martin School. Email: angelos.theodorakopoulos@oxfordmartin.ox.ac.uk

(4) School of Economics, Drexel University, USA. Email: yotov@drexel.edu

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH

Email enquiries@bankofengland.co.uk

© Bank of England 2021

ISSN 1749-9135 (on-line)

1 Introduction

Despite recent surges in bilateral trade protection, the majority of policymakers, popular observers, and leading academic economists agree that the most important obstacles to international trade take the form of country-specific policies, such as sanitary and phytosanitary measures and technical barriers to trade.¹ Many policy practitioners and academics also agree that to analyze the impact of any trade policy (either in partial or general equilibrium), one inevitably must resort to a version of the structural gravity model of trade, a.k.a. ‘the new quantitative trade model’ (Arkolakis et al., 2012; Costinot and Rodriguez-Clare, 2014).

The gravity equation, however, has been widely criticized for not being able to identify the impact of country-specific determinants of trade flows (Head and Mayer, 2014; Beverelli et al., 2018; Heid et al., 2021). More specifically, best-practices in gravity modelling have shown that the most effective way to control for general equilibrium effects and the embedding of countries into the world economy—coined by Anderson and van Wincoop (2003) as structural multilateral resistance (MR) terms—is to estimate the gravity equation with exporter(-time) and importer(-time) fixed effects (Hummels, 2001). This methodology properly controls for all country-specific characteristics and policies, however, their effect is subsumed by the fixed effects. On the other hand, not controlling for these effects with the aim of teasing out country-specific impacts on trade subjects any resulting gravity estimates to omitted variable bias (Anderson and van Wincoop, 2003).

The failure of the gravity equation to identify the impact of country-specific determinants on trade flows poses serious challenges for both partial and general equilibrium (GE) analysis. From a partial equilibrium perspective, two related challenges are noteworthy: one can neither examine the country-specific policies on international and national trade (i.e., the uniform effect) nor the full effect consisting of the uniform effect plus the discriminatory effect on international relative to national trade. This issue extends to GE analysis, as partial equilibrium estimates are often fed into models for counterfactual analysis.

To address these challenges, we derive a structural gravity equation that allows for identification of the impact of country-specific characteristics on bilateral trade flows. While our theoretical framework closely follows the standard structural gravity setting, we use it to derive a simple two-stage estimating procedure, which enables us to identify the full impact of any country-specific determinant of bilateral trade flows. Subsequently, we validate our methods by quantifying the impact of country-specific R&D expenditure on bilateral

¹See for example Lamy (2012) for policymakers’ views on the issue, as well as Kasterine (2015) and Grossman et al. (2021) for popular and academic perspectives, respectively.

trade. Of particular note, a key benefit of our estimating procedure is that it also delivers trade elasticity estimates without the need for price and/or tariff data. We implement the proposed procedure with aggregate and sectoral data, and thus a natural by-product of the analysis is that we obtain an estimate of the trade elasticity for services.

To derive our theoretical model, we capitalize and build on some of the most influential developments in the structural gravity and economic geography literatures, including Anderson (1979), Anderson and van Wincoop (2003), and Redding and Venables (2004). Moreover, consistent with the seminal work of Arkolakis et al. (2012), our gravity equation is representative of a wide class of trade models and can be derived from a series of different microeconomic foundations. For example, while in the paper we derive the model on the demand side (following Anderson, 1979 and Anderson and van Wincoop, 2003), the exact same expression can also be obtained on the supply side, *à la* Eaton and Kortum (2002). Moreover, in line with Anderson and van Wincoop (2004) and Costinot et al. (2012), our model also applies at the sector level, and can be extended to a panel setting as in Olivero and Yotov (2012), Eaton et al. (2016), and Anderson et al. (2020). The model also naturally lends itself to further extensions, some of which we discuss in the concluding section.

Our theoretical insights lead to the aforementioned simple two-stage estimating procedure, which is easy to implement with standard statistical tools and datasets. Following the latest developments in the empirical gravity literature, e.g., Head and Mayer (2014) and Yotov et al. (2016), in stage one we estimate a standard version of the structural gravity equation. The analysis at this stage delivers estimates of the vector of bilateral trade costs and MR terms, which we need to implement our structural specification in the second stage.

The second-stage analysis offers our two empirical contributions. First, the model allows for the identification of the full impact of any country-specific determinant of bilateral trade flows without inducing omitted variable bias introduced by not properly controlling for the structural MR terms (Anderson and van Wincoop, 2003; Baldwin and Taglioni, 2006). These results are robust to a number of sensitivity checks, including an instrumental variable analysis to control for potential endogeneity in our baseline estimating equation.² Second, we offer new methods to estimate trade elasticities directly from the gravity equation without the need for additional modeling or data.

Motivated by significant interest in the topic,³ we use our new methods to quantify the relationship between R&D and international trade. To this end, we rely on four measures of gross R&D expenditure (total, higher education, business enterprise, and government)

²Specifically, to address the issue of potential endogeneity of the MR terms in our second-stage analysis, we rely on and extend the methods of Feyrer (2019, 2021) and Anderson et al. (2020).

³See for example Haaland and Kind (2008), Aw et al. (2011), Peters et al. (2018), and Akcigit et al. (2018), among others.

and decompose the impact of each type of R&D spending into: (i) a uniform impact on trade regardless of whether it is domestic or international; and (ii) a discriminatory effect on international relative to domestic sales. Overall, we obtain a positive and statistically significant effect of R&D expenditure on trade, but we find that it disproportionately promotes international relative to domestic trade, thus contributing to globalization. In combination, our estimates of the uniform and the discriminatory impact imply that a 10 percent increase in total R&D expenditure translates into roughly a 3 percent increase in bilateral trade. We also document significant heterogeneity in the effects of R&D by type. Specifically, we obtain positive and significant estimates of the effect of R&D expenditure in higher education and business enterprise, but a negative estimate of the impact of government allocations.

We complement the country-specific analysis by recovering trade elasticity estimates at aggregate and sectoral levels for Manufacturing and Services, as well as Tradable versus Non-tradable goods and services sectors. A useful insight from our sectoral analysis is that we obtain an estimate of the Services trade elasticity ($\hat{\sigma}_{Services} = 7.60$), which is roughly 45 percent larger than that for Manufacturing. Moreover, we show that while the trade elasticity for Tradable sectors ($\hat{\sigma}_{Tradable} = 5.25$) is roughly in line with that for Manufacturing, that for non-tradable sectors ($\hat{\sigma}_{Non-Tradable} = 3.16$) is, unsurprisingly, much smaller by comparison.

Our paper is related to several strands of the literature. As already discussed, to derive our structural gravity equation, we capitalize and build on a number of influential studies from the theoretical trade and geography literatures.⁴ Most closely related to our paper is the seminal work of Redding and Venables (2004), who were the first to rely on gravity theory to aggregate bilateral trade costs into country-specific market access indexes to study the implications of trade for cross-country inequality. Motivated by the increased importance of country-specific trade policies and determinants of trade flows, our main contribution is to capitalize on and extend the empirical analysis of Redding and Venables (2004) into a structural estimating setting that can be used to quantify the impact of country-specific determinants of trade and to recover trade elasticities.

Our theoretical contribution advances the empirical trade literature across two important dimensions. First, we contribute to the voluminous gravity literature, which aims to quantify the effects of various determinants of trade flows by proposing methods to identify the full impact of country-specific characteristics on international trade.⁵ The work of Beverelli et al. (2018), who develop methods to identify the differential impact of country-specific characteristics on international relative to domestic trade within a structural gravity setting,

⁴We refer the reader to Anderson (2011), Arkolakis et al. (2012), Costinot and Rodriguez-Clare (2014), and Baier et al. (2018) for surveys of the theoretical gravity literature.

⁵Anderson and van Wincoop (2004), Baldwin and Taglioni (2006), Head and Mayer (2014), and Yotov et al. (2016) offer insightful surveys of the related empirical literature.

is highly related to our paper in this regard. Our model nests the specification of Beverelli et al. (2018) and, in addition, allows for identification of the uniform impact of country-specific characteristics on international and domestic trade. Thus, our setting can deliver the full (differential plus uniform) effect of any country-specific variable of interest.⁶

In implementing our methodology, we also add to the literature that studies the relationship between R&D and international trade, which, most often, analyzes these links from the firm perspective. Some excellent examples of contributions to this literature include Haaland and Kind (2008), Aw et al. (2011), Peters et al. (2018), and Akcigit et al. (2018). We complement these firm-level analyses from both a methodological and empirical angle. From a methodological perspective, we obtain estimates of the impact of R&D expenditure on trade within a theoretically-motivated multi-country empirical model. In addition to being crucial for our identification strategy, the theoretical foundation of our framework naturally lends itself to extending the analysis to study the GE implications of the relationship between R&D and trade. We view this as an interesting opportunity for future work. From an empirical and policy point of view, we quantify the total partial equilibrium impact of R&D expenditure on trade, and decompose it into uniform versus discriminatory effects. We also quantify the effects of three different categories of R&D expenditure and find significant heterogeneity by type.⁷

Second, we contribute to the literature that empirically estimates trade elasticities. Motivated by the status of the trade elasticity as one of the most crucial parameter in the international economics literature, a number of prominent studies offer methods and estimates of trade elasticities at the aggregate, sectoral, and even product level.⁸ From a methodological perspective, we add to this literature by offering a simple approach to recover trade elasticities directly from the gravity equation without the need for tariff and/or price data, and which can be applied at any level of aggregation. Given that the gravity framework is frequently used for GE analysis, a powerful implication of our methods is that they would enable researchers to recover their own trade elasticities while using the same data applied

⁶There are two other methodological differences between our study and Beverelli et al. (2018). First, their setup requires availability of data on domestic trade flows, while our methods can be applied with international data only. Second, their setting does not allow for identification of directional effects, i.e., effects on exports versus imports. Our econometric model allows for directional differences in the effects of country-specific characteristics.

⁷Importantly, however, while we demonstrate proof of concept by examining the relationship between R&D and trade, our methods easily allow to examine other country-specific trade cost determinants, such as those related to subsidies, non-discriminatory regulatory policies, technological change, etc.

⁸Without claiming to offer an exhaustive reference to the related work, we point the reader to the following influential examples: Broda and Weinstein (2006), Egger et al. (2012), Hillberry and Hummels (2013), Simonovska and Waugh (2014), Soderbery (2015), Caliendo and Parro (2015), Feenstra et al. (2018), and Fontagné et al. (2020).

in their counterfactual exercises and which correspond to the exact level of (dis-)aggregation in their GE analyses. From a practical perspective, our most valuable empirical contribution is the estimate of the trade elasticity for services.

The rest of the paper is organized as follows. In Section 2 we derive our structural gravity equation with country-specific trade costs, and translate the theoretical model into an econometric specification. The data for our analysis are described in Section 3. Section 4 brings our econometric specification to the data and presents our benchmark estimates alongside a number of sensitivity checks. After demonstrating the robustness of our benchmark estimates, we offer a proof of concept by studying the impact of R&D on bilateral trade. Section 5 demonstrates how our methodology can be used to back out trade elasticity estimates, and presents results using aggregate and sectoral data. Section 6 includes concluding remarks and directions for future work. Finally, a Supplementary Appendix includes complementary theoretical derivations, estimation results, and further data descriptions.

2 Structural Gravity with Country-specific Trade Costs

This section builds upon Anderson (1979), Eaton and Kortum (2002), Anderson and van Wincoop (2003), and Redding and Venables (2004) to derive our structural gravity equation, that simultaneously allows for the identification of the full impact of country-specific characteristics on bilateral trade flows and estimation of the trade elasticity (without the need for tariffs and/or price data). Subsequently, we translate our theoretical model into an estimating equation.

2.1 Theoretical Foundations

Consumers spend a constant share on goods from different sectors, and within sectors varieties are substitutes with a constant elasticity of substitution (CES). Hence, the preferences across sectors follow a Cobb-Douglas utility function:

$$u(C_j) = \prod_{l=1}^{\mathcal{L}} (C_j^l)^{\alpha_j^l}, \quad (1)$$

where C_j^l denotes consumption of final goods from sector l in country j , and α_j^l are the exogenously given corresponding consumption shares. Following Armington (1969), varieties in each sector are differentiated by place of origin, leading to the following CES consumer

preferences over varieties for each sector:

$$C_j^l = \left\{ \sum_{i=1}^N (\beta_i^l)^{\frac{1-\sigma^l}{\sigma^l}} (c_{ij}^l)^{\frac{\sigma^l-1}{\sigma^l}} \right\}^{\frac{\sigma^l}{\sigma^l-1}}, \quad (2)$$

where $\beta_i^l > 0$ is the CES preference parameter for goods from country i and sector l ; c_{ij}^l denotes consumption of varieties from country i of sector l in country j ; and $\sigma^l > 1$ is the elasticity of substitution between varieties in sector l . σ^l is related to the elasticity of imports with respect to variable trade costs, i.e., the trade elasticity, by $1 - \sigma^l$.

We assume a Cobb-Douglas production technology with constant returns to scale in each sector given by:

$$Y_i^l = p_i^l A_i^l (L_i^l)^{1-\gamma^l} (K_i^l)^{\gamma^l}, \quad (3)$$

where Y_i^l is the value of output in sector l in country i and p_i^l is the factory-gate price in sector l in country i . A_i^l , L_i^l and K_i^l are the corresponding levels of technology, non-tradable labor endowment, and physical capital used to produce the goods in sector l in country i , respectively.

In a trade context, one may think about other factors that could influence production, such as subsidies or non-discriminatory regulatory policies. These could easily be added to the value production function (3). Importantly, such policies will influence production and trade, but will not be identified in a state-of-the-art gravity framework. However, our two-step approach, introduced below, is able to quantify the total effect of any country-specific policy of interest in an otherwise fully theory-consistent gravity framework.

While factors are internationally immobile and can only be transferred between sectors within a country, goods of all sectors can be traded internationally. However, international trade is subject to transport costs, $t_{ij}^l \geq 1$. We follow the convenient iceberg transport cost assumption. Hence, for one unit to arrive at destination j , t_{ij}^l units must be shipped from source i . The delivered prices, p_{ij}^l , for the consumers of sector l goods from source i in destination j are therefore given by $p_{ij}^l = p_i^l t_{ij}^l$.

The budget constraint for the representative consumer is given by $\sum_l P_j^l C_j^l = \psi_j Y_j$, where ψ_j allows for exogenous trade imbalances defined as a fraction of disposable consumer income, Y_j . Y_j can be written as the sum of sectoral incomes in j , $Y_j = \sum_l Y_j^l$. Expenditure in sector l in country j is given by the constant share (i.e. the Cobb-Douglas share, α_j^l) of total disposable income, $E_j^l = \alpha_j^l \psi_j Y_j$. Equivalently, expenditure in sector l in country j can also be written as the sum over spending on all goods: $\sum_i p_{ij}^l c_{ij}^l = E_j^l$.

Utility maximization of the representative consumer leads to the demand function of

country j for goods from sector l from country i :

$$q_{ij}^l = (p_{ij}^l)^{-\sigma^l} \left(\frac{\beta_i^l}{P_j^l} \right)^{1-\sigma^l} E_j^l, \quad (4)$$

with the CES consumer price index P_j given by:

$$P_j^l = \left[\sum_{i=1}^N (\beta_i^l p_{ij}^l)^{1-\sigma^l} \right]^{\frac{1}{1-\sigma^l}} = \left[\sum_{i=1}^N (\beta_i^l p_i^l t_{ij}^l)^{1-\sigma^l} \right]^{\frac{1}{1-\sigma^l}}. \quad (5)$$

Nominal bilateral demand is obtained by multiplying both sides of (4) by the price p_{ij}^l :

$$X_{ij}^l = \left(\frac{\beta_i^l p_{ij}^l}{P_j^l} \right)^{1-\sigma^l} E_j^l. \quad (6)$$

Market clearing ensures that the sectoral production of each country is equal to the worldwide demand for its goods, $Y_i^l = \sum_{j=1}^N X_{ij}^l$. Replacing X_{ij}^l with the expression given in equation (6) and normalizing by the value of world production in sector l , Y^l , we obtain:

$$\frac{Y_i^l}{Y^l} = \sum_j \left(\frac{\beta_i^l p_{ij}^l}{P_j^l} \right)^{1-\sigma^l} \frac{E_j^l}{Y^l}. \quad (7)$$

Using $p_{ij}^l = p_i^l t_{ij}^l$ and rearranging leads to:

$$(\beta_i^l p_i^l)^{1-\sigma^l} = \frac{Y_i^l / Y^l}{\sum_j \left(\frac{t_{ij}^l}{P_j^l} \right)^{1-\sigma^l} \frac{E_j^l}{Y^l}}. \quad (8)$$

Following Anderson and van Wincoop (2003), we define the outward and inward multilateral resistance (OMR and IMR) terms, respectively, as:

$$(\Pi_i^l)^{1-\sigma^l} \equiv \sum_j \left(\frac{t_{ij}^l}{P_j^l} \right)^{1-\sigma^l} \frac{E_j^l}{Y^l}, \quad (9)$$

$$(P_j^l)^{1-\sigma^l} \equiv \sum_i \left(\frac{t_{ij}^l}{\Pi_i^l} \right)^{1-\sigma^l} \frac{Y_i^l}{Y^l}. \quad (10)$$

The OMR can be viewed as a weighted-average aggregate of all bilateral trade costs for producers in each country, while the IMR can be seen as a weighted average of all bilateral

trade costs that fall on the consumers in each country.

Using the definition of the OMR in the market-clearing equation (8) yields:

$$(\beta_i^l p_i^l)^{1-\sigma^l} = \frac{Y_i^l/Y}{(\Pi_i^l)^{1-\sigma^l}}. \quad (11)$$

We can use (11) to substitute $(\beta_i^l p_i^l)^{1-\sigma^l}$ in equation (6) to arrive at the standard structural gravity equation:

$$X_{ij}^l = \frac{Y_i^l E_j^l}{Y^l} \left(\frac{t_{ij}^l}{\Pi_i^l P_j^l} \right)^{1-\sigma^l}. \quad (12)$$

As demonstrated in Redding and Venables (2004) and Anderson et al. (2020), we can also use the expression $(\beta_i^l p_i^l)^{1-\sigma^l}$ from equation (11) to substitute factory-gate prices in the production value function (3):

$$Y_i^l = (\beta_i^l)^{-\frac{\sigma^l-1}{\sigma^l}} (A_i^l)^{\frac{\sigma^l-1}{\sigma^l}} (L_i^l)^{\frac{(\sigma^l-1)(1-\gamma)}{\sigma^l}} (K_i^l)^{\frac{(\sigma^l-1)\gamma}{\sigma^l}} (Y^l)^{\frac{1}{\sigma^l}} (\Pi_i^l)^{-\frac{\sigma^l-1}{\sigma^l}}. \quad (13)$$

A simple extension from the existing literature, which unlocks significant empirical benefits, is to exploit the structure of the gravity system even further by using equation (13) to substitute Y_i^l into equation (12). This leads to:

$$X_{ij}^l = \frac{(A_i^l)^{\frac{\sigma^l-1}{\sigma^l}} (L_i^l)^{\frac{(\sigma^l-1)(1-\gamma)}{\sigma^l}} (K_i^l)^{\frac{(\sigma^l-1)\gamma}{\sigma^l}} E_j^l (t_{ij}^l)^{1-\sigma^l}}{(\beta_i^l)^{\frac{\sigma^l-1}{\sigma^l}} (Y^l)^{\frac{\sigma^l-1}{\sigma^l}} (\Pi_i^l)^{1-\sigma^l + \frac{\sigma^l-1}{\sigma^l}} (P_j^l)^{1-\sigma^l}}. \quad (14)$$

Note that, consistent with dynamic gravity theory (Eaton et al., 2016; Anderson et al., 2020), we can add a time subscript to each of the variables in equation (14). We do this in the next subsection, where we translate our model into an estimating equation.

Before doing so, however, we reiterate that equation (14) is a structural gravity equation that capitalizes on and extends four generations of theoretical and empirical developments in the gravity literature: Anderson (1979), who derived the first gravity theory on trade; Anderson and van Wincoop (2003), who introduced the MR terms and derived the full structural gravity system; Redding and Venables (2004), who used gravity theory to aggregate bilateral trade costs into a market access index with implications for cross-country inequality; and Anderson et al. (2020), who exploited the structure of the gravity system to introduce the structural OMR as a determinant in the production function. The new step that we take by further exploiting gravity theory is to derive a structural gravity equation which enables us

to achieve two goals: (i) estimate the full effects of country-specific characteristics on trade, and (ii) obtain estimates of the trade elasticity directly from trade data, i.e., without the need for data on prices and/or tariffs.

2.2 From Theory to Empirics

In what follows, we rely on and extend the latest developments in the empirical gravity literature to translate the theoretical gravity equation (14) into an estimating model.⁹ However, as noted earlier and also consistent with theory, we introduce a time subscript to our specification. Following the recommendations of Santos Silva and Tenreyro (2006, 2011) we use the Poisson pseudo-maximum likelihood (PPML) estimator, which accounts for potential heteroscedasticity and zero trade flows in trade data,¹⁰ to obtain our main estimates. As such, our estimating equation becomes:

$$X_{ij,t} = \exp[\alpha_1 \ln(A_{i,t}) + \alpha_2 \ln(L_{i,t}) + \alpha_3 \ln(K_{i,t}) + \alpha_4 \ln(E_{j,t}) + \alpha_5 \ln(t_{ij,t})] \times \exp[\alpha_6 \ln(\Pi_{i,t}^{1-\sigma}) + \alpha_7 \ln(P_{j,t}^{1-\sigma}) + \alpha_8 \beta_i + \alpha_9 Y_t] \times \epsilon_{ij,t}, \quad (15)$$

where the structural interpretation of the coefficients are as follows: $\alpha_1 = (\sigma - 1)/\sigma$; $\alpha_2 = (\sigma - 1)(1 - \gamma)/\sigma$; $\alpha_3 = (\sigma - 1)\gamma/\sigma$; $\alpha_4 = 1$, $\alpha_5 = 1 - \sigma$; $\alpha_6 = (1 - \sigma)/\sigma$;¹¹ $\alpha_7 = -1$; and $\alpha_8 = \alpha_9 = (1 - \sigma)/\sigma$.

As underscored in Anderson and van Wincoop (2003), a major challenge with estimating (15) is that the MR terms are not observable. And, as mentioned above, the standard practice of using exporter-time and importer-time fixed effects will subsume any country-specific policies or characteristics, rendering their identification infeasible.

To address this issue, we propose a simple two-step procedure. In Step 1, we estimate the following standard version of equation (15) with exporter-time and importer-time fixed effects:

$$X_{ij,t} = \exp[\alpha_5 \ln(t_{ij,t}) + \phi_{i,t} + \psi_{j,t}] \times \epsilon_{ij,t}. \quad (16)$$

Next, we capitalize on the additive property of PPML (Arvis and Shepherd, 2013; Fally, 2015) to recover the power transformations of the MRs from the estimates of the fixed

⁹We omit the sector superscript, since our model is estimated with aggregate data.

¹⁰Note that, at the level of aggregation in this analysis, we do not observe any zeros.

¹¹As we recover the MRs to the power of $1 - \sigma$ from the fixed effects, we transform the power of the corresponding variable as follows: $1 - \sigma + (\sigma - 1)/\sigma = (2\sigma - \sigma^2 - 1)/\sigma = (1 - \sigma)[(2\sigma - \sigma^2 - 1)]/[(1 - \sigma)\sigma]$, which leads to $\alpha_6 = -(2\sigma - \sigma^2 - 1)/[(1 - \sigma)\sigma] = (1 - \sigma)^2/[(1 - \sigma)\sigma] = (1 - \sigma)/\sigma$.

effects in equation (16) in combination with data on output and expenditure:

$$\widehat{\Pi}_{i,t}^{1-\sigma} = \frac{Y_{i,t}}{\exp(\widehat{\phi}_{i,t})} \times \frac{E_{0,t}}{Y_t}, \quad (17)$$

and

$$\widehat{P}_{j,t}^{1-\sigma} = \frac{E_{j,t}}{\exp(\widehat{\psi}_{j,t})} \times \frac{1}{E_{0,t}}, \quad (18)$$

where $E_{0,t}$ is the expenditure of the country that has been selected as a numeraire, i.e., whose importer fixed effects are dropped when PPML is estimated without a constant.¹²

In step 2 we replace the MR estimates from above into equation (15) to obtain:

$$\begin{aligned} X_{ij,t} &= \exp[\alpha_1 \ln(A_{i,t}) + \alpha_2 \ln(L_{i,t}) + \alpha_3 \ln(K_{i,t}) + \alpha_4 \ln(E_{j,t}) + \alpha_5 \ln(t_{ij,t})] \times \\ &\times \exp[\alpha_6 \ln(\widehat{\Pi}_{i,t}^{1-\sigma}) + \alpha_7 \ln(\widehat{P}_{j,t}^{1-\sigma}) + \alpha_8 \beta_i + \alpha_9 Y_t] \times \epsilon_{ij,t}. \end{aligned} \quad (19)$$

Several notes are in order. First, we include exporter fixed effects, ϕ_i , to control for the exporter-specific preference parameters in our model, β_i , as well as year fixed effects, ϕ_t , to control for world output, Y_t . In addition, these fixed effects will absorb any other observable and unobservable characteristics along these dimensions. Second, from theory, the estimate of the coefficient on expenditure (α_4) and on the IMR (α_7) should be equal to one. Note, however, there is not much to be gained from the inclusion of these variables in the current setting. Therefore, in addition to regressions where we include expenditure and IMRs, we also estimate specifications where we include importer-time fixed effects and specifications where we use the importer-fixed effects obtained from the first step as constraints.¹³

Third, we have to approximate for trade costs $t_{ij,t}$. To do so, we take two different approaches. First, we approximate bilateral trade costs very flexibly with symmetric bilateral fixed effects (μ_{ij}), symmetric bilateral linear trends ($\mu_{ij} \times t$), and time-varying international border effects ($BRDR_{ij} \times \phi_t$, where $BRDR_{ij}$ is a dummy variable equal to one if countries i and j are separated by an international border and zero otherwise). Further, we also include three controls for common membership in regional trade agreements (RTAs), the European

¹²See Fally (2015) and Anderson et al. (2018) for further details on the calculation of the structural MRs from the estimates of PPML fixed effects. Two other methods can be used to deliver exactly the same estimates of the power transformations of the multilateral resistances. First, we could implement the iterative procedure of Anderson and van Wincoop (2003). Second, which is standard in the literature, we can solve the theoretical system for the structural MR terms. We capitalize on the PPML fixed effects properties for our two-step approach as it is the simplest possible way to obtain the MRs and, therefore, makes our procedure straightforward to implement.

¹³Note that we also could unpack the importer fixed effects, as we unpack the exporter fixed effects using the production function. This would allow us to investigate the drivers of the importer fixed effects and to investigate country-specific effects on the importer and on the exporter side, for example. We see this as a fruitful area for future research.

Union (EU), and the World Trade Organization (WTO). Overall, this specification captures any bilateral specific components, trend in bilateral changes, common effects of any global trends (e.g., in communication, transportation, technology, etc.) and specific trade policy effects via RTAs and EU or WTO membership as trade costs.¹⁴ As a second alternative approach, we use the trade cost estimates from the first step, constructed using the exact same fixed effects and variables, and constrain bilateral trade costs with these estimates.

Fourth, and finally, we have to control for $A_{i,t}$. To do so, we use a measure of total factor productivity (TFP) as a proxy. In addition, the employed fixed effects will, at least partially, mitigate the known limitation of any TFP measure, namely that it is an estimate of an unobservable term and thus might not perfectly capture true productivity differences across countries and over time. More specifically, the exporter fixed effects will account for any country-specific productivity differences, while the (importer)-time fixed effects will control for global technology/productivity trends and changes over time.¹⁵

To sum up, when accounting for all of the above considerations, our main estimating equation is expressed as:

$$\begin{aligned}
X_{ij,t} &= \exp[\alpha_1 \ln(A_{i,t}) + \alpha_2 \ln(L_{i,t}) + \alpha_3 \ln(K_{i,t}) + \alpha_4 RTA_{ij,t} + \alpha_5 EU_{ij,t}] \times \\
&\times \exp[\alpha_6 WTO_{ij,t} + \alpha_7 \ln(\widehat{\Pi_{i,t}^{1-\sigma}}) + \phi_i + \psi_{j,t} + \mu_{ij} + \mu_{ij} \times t] \times \\
&\times \exp[BRDR_{ij} \times \phi_t] \times \epsilon_{ij,t}.
\end{aligned} \tag{20}$$

Note that the importer-time fixed effects also absorb overall time-specific fixed effects. Importantly, from the obtained coefficient for the effect of the OMR (α_7) we can back out an estimate of σ .¹⁶

Two final considerations should be taken into account when estimating equation (20). First, we cluster standard errors by country-pair to allow for correlation in the error term at this level. Second, the MRs are potentially endogenous by construction. To address this issue, we capitalize on and extend the methods of Feyrer (2019), Feyrer (2021), and Anderson et al. (2020) by constructing a series of instruments for the OMRs.¹⁷

¹⁴Note that, due to the use of domestic trade flows, we also could add country-specific variables, like institutions or non-discriminatory trade policies, if we interact them with the international border dummy, as in Beverelli et al. (2018). We demonstrate this in Section 4.3, where we apply our methods to study the impact of country-specific R&D expenditure on international trade.

¹⁵Alternatively, for robustness we assume a first-order autoregressive process for productivity and use the structure of our model to replace technology by explanatory variables. Details are provided in Supplementary Appendix A.

¹⁶Specifically, this is $\sigma = 1/(1 + \alpha_7)$. As we discuss in section 5, however, the most theory-consistent approach to back out σ is to estimate equation (20) while imposing all available constraints on the structural parameters suggested by the model.

¹⁷See section 4.2 for further detail.

Equation (20) has two main advantages in relation to standard gravity equations. Consistent with our main objective, it enables us to obtain estimates of country-specific determinants of trade flows and to decompose their effects into uniform versus discriminatory impacts on international relative to internal trade. We demonstrate this in Section 4.3, where we apply our methods to study the impact of country-specific R&D expenditure on international trade. However, the estimating equation can accommodate any country-specific variable. Further, equation (20) enables us to recover (sectoral) estimates of the trade elasticity directly from bilateral trade data. We demonstrate this in Section 5, where we present trade elasticity estimates for Aggregate, Manufacturing, and Services trade, as well as for Tradable and Non-tradable sectors.

3 Data Description and Sources

To conduct our empirical analysis, we combine data from several sources. First, we source data on trade, output, and expenditure from the World Input-Output Tables (WIOTs), one of two datasets made publicly available by the World Input-Output Database (WIOD). WIOTs are available for 43 countries,¹⁸ cover years 2000-2014, and provide consistent and harmonized information, in USD, on bilateral (intermediate and final) international and domestic trade flows for 56 disaggregate goods and services industries (see Appendix Table B1 for a detailed breakdown).¹⁹

Next, we match the WIOTs for each year with country-specific data on factors of production (capital and labor compensation) from the WIOD Socio Economic Accounts (SEA), which cover the same countries, time span, and sector disaggregation. To ensure consistency across datasets, we convert the capital and labor compensation variables, available in national currency, to USD using the same exchange rates used in the production of the WIOTs. In addition, before summing the data to our desired level of aggregation, we exclude disaggregate WIOD industries for which we observe positive total trade flows across all destination countries within a given exporter-sector-year, but which do not exhibit positive values for capital or labor compensation at this same level. Overall, this affects a trivial share of the data (4.6 percent).

The main advantages of using WIOD for our analysis are threefold. First, the WIOTs

¹⁸See Table B2 for a full list of countries included in our analysis. Beyond the countries listed in Table B2, the WIOTs contain information on a ‘rest of world’ (RoW) aggregate and Taiwan. We exclude Taiwan from our analysis due to limited coverage of capital compensation. We exclude the RoW from our analysis due to paucity of data on various aggregates such as bilateral trade agreements at the world level.

¹⁹At this stage, we also clean the data by setting negative trade flow values to zero, if and when they occur. This affects only final goods/services trade, and represents only 0.18 percent of the data.

allow for consistency with our structural model insofar as they include information on both domestic and international trade flows. Second, given the wide number of sectors available, we are able to obtain disaggregate results and provide sector-specific trade elasticities. This is particularly important for services, and is a novel contribution of our analysis. Finally, combining information from WIOT and SEA allows us to obtain data on the country-specific factors of production at a disaggregate level, which is important for our analysis at the sectoral level.

Subsequently, we match all WIOD data at the aggregate level with information on TFP at current PPPs and population from version 10.0 of the Penn World Table (PWT), dummy variables for common bilateral EU and WTO membership from CEPII, and bilateral RTA membership from the Regional Trade Agreements Database from Egger and Larch (2008). Finally, we source country-specific data on gross domestic expenditure on R&D at current PPPs per capita from the Organisation for Economic Co-Operation and Development Main Science and Technology Indicators Database. We include four variables in our country-specific analysis: total gross domestic expenditure on R&D at current PPPs per capita, which is then broken down into three sub-categories for higher education expenditure, business enterprise expenditure, and total government allocations.

4 Estimating Country-specific Trade Costs

We start this section by applying our procedure to aggregate data. We subsequently perform a series of sensitivity checks and demonstrate that our methods and results are robust to various endogeneity concerns. In this vein, we undertake three separate instrumental variable (IV) approaches to correct for potential endogeneity in our baseline estimating equation. Finally, upon demonstrating the strength of our methodology and robustness of our empirical results, we provide proof of concept by examining the links between country-specific R&D expenditure and bilateral trade.

4.1 Benchmark Estimates

Table 1 presents the empirical results for our baseline specification (20). The two panels of Table 1 are based on different approaches to proxy bilateral trade costs. Estimates in panel A are obtained with fully flexible bilateral trade costs. As described in section 2.2, we approximate bilateral trade costs by controlling for symmetric bilateral fixed effects, bilateral linear trends, time-varying international border effects, and dummy variables for common EU, WTO, and RTA membership. The estimates in panel B are obtained after

constraining the vector of bilateral trade costs in the second stage analysis to be the vector of corresponding estimates that we obtain from the first stage.

Column (1) presents our most flexible specification, whereby we include importer-time fixed effects directly. We note that the estimates of the effects of RTA, EU, and WTO membership exhibit expected signs and magnitudes in line with the literature. Further, the point estimates on labor and capital fall within the theoretical bounds of zero to one and, with a ratio of roughly 2/3 and 1/3, respectively, have plausible magnitudes. Our point estimate of the power transformation of the OMR term, $\ln(\widehat{\Pi_{i,t}^{1-\sigma}})$, is equal to -0.857 and is statistically significant at the 1% level.

Table 1: Baseline Results with Country-level Data

	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A			Panel B		
	Flexible bilateral trade costs with constrained importer-time FEs:			Constrained bilateral trade costs with constrained importer-time FEs:		
	None	Partial	Full	None	Partial	Full
$\ln(A_{i,t})$	0.070* (0.042)	0.085* (0.046)	0.085* (0.048)	0.046 (0.093)	0.053 (0.049)	0.053 (0.050)
$\ln(L_{i,t})$	0.579*** (0.054)	0.561*** (0.032)	0.579*** (0.028)	0.555*** (0.072)	0.552*** (0.039)	0.563*** (0.032)
$\ln(K_{i,t})$	0.339*** (0.053)	0.346*** (0.040)	0.361*** (0.040)	0.382*** (0.073)	0.383*** (0.048)	0.392*** (0.045)
$\ln(\widehat{\Pi_{i,t}^{1-\sigma}})$	-0.857*** (0.062)	-0.817*** (0.074)	-0.877*** (0.047)	-0.842*** (0.079)	-0.833*** (0.097)	-0.878*** (0.054)
$RTA_{ij,t}$	0.072** (0.035)	0.077** (0.036)	0.077** (0.036)			
$EU_{ij,t}$	0.032 (0.034)	0.038 (0.039)	0.015 (0.032)			
$WTO_{ij,t}$	0.051 (0.058)	0.043 (0.057)	0.042 (0.057)			
$\ln(E_{j,t})$		0.969*** (0.034)			0.972*** (0.046)	
$\ln(\widehat{P_{j,t}^{1-\sigma}})$		-0.930*** (0.071)			-0.949*** (0.095)	
Observations	26,460	26,460	26,460	26,460	26,460	26,460

Notes: This table presents results for equation (20) using country-level data following the two-stage approach. Estimates in panel A are obtained with fully flexible bilateral trade costs. Column (1) includes importer-time FEs directly. Column (2) includes expenditure and estimates of the power transformation of the IMR term from equation (18) instead of importer-time fixed effects. Further, column (3) includes the computed vector of importer-time FEs from the first stage, with its parameter constrained to unity. Estimates in panel B are obtained after constraining the vector of bilateral trade costs to be the vector of corresponding estimates that we obtain from the first stage. Results in columns (4)-(6) are obtained following the same steps for columns (1)-(3). Column (6) represents our preferred baseline specification. Standard errors clustered by country-pair in parentheses. * $p < 0.10$, ** $p < .05$, *** $p < .01$. See text for further details.

In the next two columns of panel A, we gradually constrain the specification from column (1) given estimates from the first stage and guided by theory. Specifically, we vary the degree to which we constrain the importer-time fixed effects. In column (2), we include expenditure and estimates of the power transformation of the IMR term from the first stage, corresponding to equation (18) instead of importer-time fixed effects. The estimates on both new variables are statistically significant and with signs as predicted by theory. Moreover, in line with our priors based on theory, we fail to reject the null hypothesis that the estimates for expenditure are statistically different than one. Importantly, the estimates on all other variables are very close to those in column (1). Next, in column (3), we directly include the computed vector of importer-time fixed effects from the first stage, with its parameter constrained to unity. Once again, results remain comparable to those from column (1).

The estimates in panel B are obtained after constraining the vector of bilateral trade costs in the second-stage estimation to be the vector of corresponding estimates that we obtain from the first stage. Results in columns (4)-(6) are obtained following exactly the same steps for the specification of the importer-time fixed effects that we took to in columns (1)-(3). Specifically, in column (4) we include importer-time fixed effects. Then, in column (5), we include the structural covariates for expenditure and the IMR term from the first stage. Finally, in column (6), we constrain the estimates on the two structural variables to unity. Column (6) represents our preferred specification as it aligns most closely with theory. Overall, however, the estimates in column (6) are as expected, and are strikingly similar to all other results in Table 1.

4.2 Endogeneity Concerns and Robustness Analysis

There are at least three potential reasons for endogeneity concerns with the key structural OMR covariate, $\ln(\Pi_{i,t}^{1-\sigma})$, in equation (20). First, it is endogenous by construction because it includes own-country national income. Second, in addition to the direct link to own size, it could be endogenous due to indirect links between own national income and the national incomes of all other countries. Finally, there could also be an indirect relationship between own national income and the IMRs of all other countries. To mitigate these concerns, we follow and extend the existing literature to construct three different IVs. In what follows, we describe each IV and discuss the corresponding estimates in turn.²⁰

²⁰We rely on the generalized method of moments estimation technique for panel count data models with endogenous regressors, which are discussed in detail in Windmeijer and Santos Silva (1997) and Windmeijer (2008). As discussed in this literature, when controlling for unobserved heterogeneity, panel count data models might suffer from incidental parameter problems with inconsistent estimates unless the time dimension is large. Note, however, that this is not of concern in our estimation setting since our variable of interest varies at the exporter-importer-year (ijt) level, and thus we have a relatively large number of observations

Table 2 presents results, where, to ease comparison, column (1) reproduces results from our preferred specification from column (6) of Table 1. In our first IV, IV-1, we follow Anderson et al. (2020) and re-construct the OMR term using base-year weights. Mechanically, we compute: $\Pi_{i,t}^{1-\sigma} = \sum_{j \neq i} \left(\frac{t_{ij,t}}{P_{j,t}} \right)^{1-\sigma} \frac{W_{j,2000}}{W_{2000}}$, where W represents base-year weights either using number of employees from SEA or base-year population from PWT.²¹ Thus, we simultaneously eliminate the direct own-size links by construction as well as residual endogeneity due to links between national income and incomes of other countries. Results under this IV strategy, when using the number of employees as weights,²² are presented in column (2). We note that they are very similar to our baseline estimates in column (1).

Table 2: Baseline Results with IVs for OMR Terms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		No. Employees Weights			Population Weights		
	Baseline	IV-1	IV-2	IV-3	IV-1	IV-2	IV-3
$\ln(A_{i,t})$	0.053 (0.050)	0.053 (0.051)	0.053 (0.051)	0.053 (0.050)	0.053 (0.050)	0.053 (0.050)	0.053 (0.050)
$\ln(L_{i,t})$	0.563*** (0.032)	0.585*** (0.103)	0.584** (0.231)	0.569*** (0.119)	0.547*** (0.106)	0.543** (0.256)	0.569*** (0.119)
$\ln(K_{i,t})$	0.392*** (0.045)	0.391*** (0.046)	0.391*** (0.044)	0.392*** (0.044)	0.394*** (0.046)	0.394*** (0.046)	0.392*** (0.044)
$\ln(\widehat{\Pi_{i,t}^{1-\sigma}})$	-0.878*** (0.054)	-0.920*** (0.195)	-0.918** (0.441)	-0.889*** (0.216)	-0.848*** (0.194)	-0.842* (0.482)	-0.889*** (0.216)
Observations	26,460	26,460	26,460	26,460	26,460	26,460	26,460

Notes: This table presents results using country-level data. Column (1) presents baseline results for equation (20). In all subsequent columns we correct for potential endogeneity from the OMR term in equation (20) through different instruments. Column (2) solves for IMR and OMR terms using base-year weights and $t_{ij,t}$ estimated from the first stage. Column (3) solves for IMR and OMR terms using base-year weights and $t_{ij,t}$ estimated following Feyrer (2019). Column (4) uses a base-year weighted average of t_{ij} estimated following Feyrer (2019). Columns (2)-(4) use number of employees in year 2000 for base-year weights. Columns (5)-(7) correct for potential endogeneity in equation (20) following the same procedure as columns (2)-(4), but instead use population in year 2000 for base-year weights. All IV columns exclude within-country values. Standard errors clustered by country-pair in parentheses. * $p < 0.10$, ** $p < .05$, *** $p < .01$.

Our second IV, IV-2, extends the methods in IV-1 to account for potential additional higher-order endogeneity introduced by the bilateral fixed effects in our trade cost vector. Following the recommendations of Feyrer (2019), we build upon the approach taken in IV-1 but now compute $\widetilde{t_{ij,t}}$ from a first-stage regression where we regress international bilateral gross exports on time-varying distance ($dist_{ij,t} \equiv dist_{ij} \times \phi_t$) and use this trade cost vector to identify the exporter (i) and time (t) fixed effects.

²¹Specifically, given the trade cost vector and the base-year weights, we solve for the system of equations that come from the OMR and IMR terms, the latter being $P_{j,t}^{1-\sigma} = \sum_{i \neq j} \left(\frac{t_{ij,t}}{\Pi_{i,t}} \right)^{1-\sigma} \frac{W_{i,2000}}{W_{2000}}$.

²²Below we also report and discuss estimates that use population weights.

to re-compute MR terms. As such, we compute $\Pi_{i,t}^{1-\sigma} = \sum_{j \neq i} \left(\frac{\widetilde{t}_{ij,t}}{P_{j,t}} \right)^{1-\sigma} \frac{W_{j,2000}}{W_{2000}}$, and present results using employee weights in column (3). Results remain in line with our baseline estimates.

Finally, our third IV, IV-3, uses the recomputed trade cost vector from IV-2 and takes a weighted average for each exporter across all importers, using base-year values as weights. Otherwise stated, we compute $\Pi_{i,t}^{1-\sigma} = \sum_{j \neq i} \left(\widetilde{t}_{ij,t} \right)^{1-\sigma} \frac{W_{j,2000}}{W_{2000}}$ and present results when using employee weights in column (4). Results remain similar to our baseline specification.²³

Columns (5)-(7) of Table 2 follow the same estimation approach as columns (2)-(4), except employ population weights. Results are very similar under the two estimation approaches, albeit with slightly less negative estimates on the OMR terms when using population weights.

To further test the robustness of our results, we conduct several additional experiments, presented in Table 3, where once again, column (1) reproduces our baseline estimates from column (6) of Table 1. To motivate our first robustness check, recall that the first-stage estimation allows us to directly retrieve the error term, $\epsilon_{ij,t}$. We take advantage of this information in column (2), where we use total exports corrected for the error term from the first stage estimation as a dependent variable. Comparing with our baseline results in column (1), point estimates and standard errors are virtually identical.

In columns (3) and (4), we exploit the fact that the SEA database provides additional measures of capital and labor, beyond what we include in our baseline estimation. As such, in column (3), we replace capital compensation with a measure of nominal capital stock. While this yields a smaller point estimate overall for $K_{i,t}$, the point estimate on the OMR term is robust and in line with baseline results. Similarly, in column (4), we replace our measure of labor compensation with compensation of employees.²⁴ This yields a slightly higher estimate of $L_{i,t}$ but, as was the case for column (3), the point estimate on the OMR term remain robust.

In column (5) we implement three-way clustering of standard errors by exporter, importer, and year to allow for correlation of the error term across all three dimensions. Standard errors are slightly larger, with the overall significance of our baseline results unchanged. Finally, to conclude the robustness analysis, in column (6) we employ an alternative method to control for productivity. This is a potentially important improvement because the TFP term used so far is based on estimates of a production function which does not necessarily align with all underlying assumptions in our structural framework. Thus, we aim to show

²³A possible explanation as to why IV-2 and IV-3 deliver larger standard errors for the OMR term as compared to our baseline and IV-1 is that since our panel is relatively short, spanning years 2000-2014, potentially we do not capture large amounts of variation over time when computing $dist_{ij,t}$.

²⁴Note that labor compensation is potentially a more comprehensive measure, insofar as it includes compensation to all labor types, e.g. the self-employed.

that our results are not sensitive to the productivity measure at hand. Accordingly, we no longer rely on estimated TFP from PWT (or any other source). Instead, we assume that productivity follows a Markov (first order auto-regressive) process, which allows us to estimate a version of equation (20) without the need to first estimate TFP (see Supplementary Appendix A for a detailed description of the methodology). This approach is in the same spirit as both dynamic panel (Blundell and Bond, 2000) and proxy variable methods (Olley and Pakes, 1996) for estimating production functions. Results in column (4) reveal that the new approach to control for productivity leads to changes in the levels and in the relative shares of labor and capital, which remain plausible. More importantly for our purposes, the estimate on the OMR remain very similar to our previous results, lending support to this approach overall.

Table 3: Robustness

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline results	Baseline hat	Alternate measures of: $K_{i,t}$	$L_{i,t}$	Three-way cluster	Alternative approach
$\ln(A_{i,t})$	0.053 (0.050)	0.053 (0.051)	0.191*** (0.070)	0.089** (0.043)	0.053 (0.051)	
$\ln(L_{i,t})$	0.563*** (0.032)	0.563*** (0.033)	0.719*** (0.055)	0.571*** (0.035)	0.563*** (0.035)	0.438*** (0.034)
$\ln(K_{i,t})$	0.392*** (0.045)	0.392*** (0.049)	0.142** (0.065)	0.369*** (0.046)	0.392*** (0.050)	0.422*** (0.025)
$\ln(\widehat{\Pi_{i,t}^{1-\sigma}})$	-0.878*** (0.054)	-0.878*** (0.055)	-0.884*** (0.047)	-0.862*** (0.056)	-0.878*** (0.062)	-0.847*** (0.037)
Observations	26,460	26,460	26,460	26,460	26,460	24,696

Notes: This table presents results for equation (20) using country-level data in column (1) following the two-stage approach. Column (2) uses gross exports corrected for the error term from the first-stage regression as a dependent variable. Column (3) constrains the coefficients to ensure constant returns to scale. Column (4) uses an alternative estimation approach which assumes a first order auto-regressive process for productivity, as discussed in Supplementary Appendix A. Columns (5) and (6) use nominal capital stock and compensation of employees from SEA as alternate measures of capital and labor, respectively. Column (7) implements three-way clustering at the i j and t level. * $p < 0.10$, ** $p < .05$, *** $p < .01$. See text for further details.

4.3 Proof of Concept: Links between R&D and Trade

Having demonstrated the robustness of our baseline estimating equation, we now show how the structural gravity model can deliver estimates of the full effects of country-specific determinants of international trade, focusing on the impact of R&D expenditure as an example.²⁵ In doing so, we decompose the impact of R&D into two components: (i) a uniform impact

²⁵As noted previously, however, our methods and analysis apply to any country-specific variable.

on sales regardless of destination, and (ii) a discriminatory effect on international relative to domestic sales. We offer novel estimates by exploring four R&D measures: (i) total gross expenditure on R&D at current PPPs per capita (Total); which is then broken down into sub-categories for (ii) higher education (H. Edu.); (iii) business enterprise (Business); and (iv) government allocations (Govt.).²⁶ Finally, we show that our methods deliver results that can be very different than those produced when applying alternative procedures that the literature has proposed to obtain estimates of the effects of country-specific variables.

Table 4 presents our findings, where columns (1)-(4) offer results for each of the four R&D variables, respectively. The four panels, A-D, report estimates from alternative approaches to obtain estimates of R&D effects on bilateral trade flows.

Panel A follows the two-stage estimation procedure as detailed throughout the paper. Note that in the first-stage estimating equation we now include the interaction of $R\&D_{i,t}$ with international borders, as in Heid et al. (2021) and Beverelli et al. (2018). This allows us to obtain an estimate of the partial effect of R&D expenditure on international trade relative to domestic sales (i.e. the interaction term $\ln(R\&D_{i,t}) \times BRDR_{ij}$). Overall, we find a positive and statistically significant effect of R&D expenditure on international relative to domestic trade, suggesting that innovation in R&D promotes international trade disproportionately, thus contributing to globalization. However, these effects vary by R&D type. Specifically, we obtain positive and significant estimates of the impact of higher education and business enterprise R&D expenditure (columns 2 and 3, respectively), but a negative, sizable, and statistically significant estimate of the impact of government allocations (column 4). One possible explanation for the latter result is ‘home bias’ in government spending. As such, we find this result intuitive but also interesting and provoking.

We next incorporate $R\&D_{i,t}$ into our second-stage estimation, i.e. by adding an extra term into our baseline estimating equation (corresponding to Table 1, column 6). The estimates in ‘Stage 2’ of panel A reveal that the impact of each type of R&D expenditure on sales is positive and statistically significant. This is another intuitive finding, as we would expect that innovation in R&D should promote efficiency. Furthermore, we note that the estimates of the effects of the different types of R&D expenditures are similar in magnitude.

In combination, the estimates from ‘Stage 1’ and ‘Stage 2’ of panel A enable us to obtain the total impact of each type of R&D on international trade; we report these indexes in the last row of panel A. Two main findings stand out. First, the total impact of R&D expenditure in higher education is the strongest, closely followed by the effects of business

²⁶As discussed earlier, by quantifying the impact of R&D on trade with our methods we complement a series of excellent studies, e.g., Haaland and Kind (2008), Aw et al. (2011), Peters et al. (2018), and Akcigit et al. (2018), which study this relationship from the perspective of the firm.

enterprise expenditure. Specifically, our estimates suggest that a 10 percent increase in R&D expenditure in higher education translates into about a 2 percent increase in international trade, while the corresponding estimate for business enterprise expenditure is 1.64 percent. Second, while the positive uniform effect of government R&D allocations work in the opposite direction of its negative discriminatory impact on international trade, the former is not enough to outweigh the latter. Therefore, according to our calculations, the net impact of government R&D expenditures on international trade is quite small, but still negative.

Table 4: Country-specific Results - R&D

	(1) Total	(2) H. Edu.	(3) Business	(4) Govt.
Panel A: Stage 1				
$\ln(R\&D_{i,t}) \times BRDR_{ij}$	0.244*** (0.094)	0.115* (0.059)	0.113* (0.060)	-0.176*** (0.049)
Stage 2				
$\ln(R\&D_{i,t})$	0.081*** (0.014)	0.075*** (0.019)	0.051*** (0.013)	0.082*** (0.018)
Total effect	0.325	0.191	0.164	-0.094
Panel B: Naïve gravity				
$\ln(R\&D_{i,t}) \times BRDR_{ij}$	0.232** (0.095)	0.114** (0.058)	0.104 (0.072)	-0.139** (0.058)
$\ln(R\&D_{i,t})$	0.082 (0.055)	0.016 (0.032)	0.039 (0.041)	0.069* (0.038)
Total effect	0.314	0.131	0.143	-0.070
Panel C: FE PPML, Stage 2				
$\ln(R\&D_{i,t})$	0.069*** (0.021)	0.046*** (0.017)	0.031* (0.019)	0.061*** (0.021)
Total effect	0.313	0.161	0.144	-0.115
Panel D: FE OLS, Stage 2				
$\ln(R\&D_{i,t})$	0.014 (0.025)	0.037** (0.014)	-0.003 (0.014)	0.059** (0.027)
Total effect	0.258	0.152	0.110	-0.117

Notes: The country-specific variable we include is gross domestic expenditure on R&D at current PPPs per capita (column 1), broken down into sub categories for: higher education (H. Edu., column 2); business enterprise expenditure (Business, column 3); and total government allocations (Govt., column 4). For brevity, only country-specific variables and their interactions with border dummy are included. See Supplementary Appendix C for full results. Standard errors clustered by country-pair in parentheses. * $p < 0.10$, ** $p < .05$, *** $p < .01$. See text for further details.

We next break from our two-stage procedure and experiment with alternative estimation techniques of country-specific variables from the literature. In this vein, Panel B presents results from a “naïve” one-stage gravity specification, whereby we exclude the exporter-time fixed effect such that we can include time-varying country-specific characteristics. As noted above, this estimation procedure does not properly account for the structural OMR of Anderson and van Wincoop (2003). As such, results are bound to be biased for this reason; also known as the ‘gold medal mistake’ critique of Baldwin and Taglioni (2006). Nonetheless, we explore this naïve setup since it has been frequently employed in the literature to obtain estimates of various country-specific characteristics.

Concretely, the econometric model that we employ includes all bilateral covariates that are used in our main specification (i.e., country-pair fixed effects interacted with a time trend, time-varying border variables, and controls for common RTA, EU and WTO membership) as well as the full set of importer-time fixed effects which fully accounts for IMRs. However, instead of using exporter-time fixed effects, we naïvely include exporter fixed effects, year fixed effects, GDP as a proxy for exporter size, and an a-theoretical remoteness index, which, following the vast gravity literature, we construct as importer GDP-weighted bilateral distance. Consistent with our main specification, and due to the use of domestic trade flows, we are able to obtain estimates of the differential effect of R&D on international relative to domestic trade (through the interaction of $\ln(R\&D_{i,t}) \times BRDR_{ij}$) as well as its uniform impact by directly including the $\ln(R\&D_{i,t})$ covariate.

We draw the following conclusions based on panel B results. First, the estimates of the differential effect of R&D on trade are similar to the corresponding indexes in panel A, however, the estimate of the impact of business enterprise expenditure is no longer statistically significant. Second, all estimates of the uniform effects of R&D (except for that of total R&D) are smaller in magnitude and, more importantly, only one of them (on government allocations) is statistically significant. Thus, we do not find evidence of any uniform effect of R&D on aggregate sales. This is in contrast with expectations as well as main findings from panel A, thus casting doubt on this specification. Based on this, we conclude that the naïve one-stage gravity specification may result in significant biases in the country-specific estimates of interest. Furthermore, this approach cannot deliver estimates of trade elasticities, a potentially significant drawback.

The specification that we employ in panel C aims at comparing our two-stage procedure with the common practice to regress fixed effects recovered from the first-stage regression on explanatory variables (including country-specific variables) in a second-step. In order to do so, motivated by our theory, we implement the following two-stage procedure. The first-stage estimating equation is identical to that in Panel A. As before, this equation will deliver

estimates of the differential effects of R&D on international relative to domestic trade. Then, the second-stage estimating equation uses the fitted values of the exporter-time fixed effects from the first stage and regresses them on all exporter-specific variables in our second-stage specification from panel A, including R&D expenditure. As the exporter-time fixed effects are a multiplicative function of output and OMRs (see equation (17)), we follow the same logic as for the trade equation and estimate the second-stage using PPML.

Based on the new estimates, we conclude that the procedures we employ in panels A and C deliver similar estimates. Since the two procedures are closely related by design, we found the similarity between the estimates in panels A and C reassuring. The implication is that, as long as the structural properties of the gravity model are preserved, researchers may select their preferred method. We prefer the model in Panel A because it uses bilateral trade flows (our outcome of interest) as the dependent variable in the second stage in a theoretically founded way. Hence, it does not rely on fixed effects to simultaneously decompose the effects of R&D and to estimate its full impact. Moreover, given that Panel C relies upon fixed effects as a dependent variable, assumptions about the pattern of heteroscedasticity and error term variance inherently deviate from those in Panel A. Overall, the results in panel C reinforce our main contribution, which is to demonstrate that it is possible to obtain estimates of the full impact of country-specific variables within a properly specified structural gravity model.

We conclude with an analysis of results in panel D, where we employ the same two-stage procedure used in panel C, however, instead of using the PPML estimator in stage-two, we rely on OLS. The use of OLS for the second-stage analysis is consistent with the standard two-stage approach from the related literature (Head and Mayer, 2014). Interestingly, the estimated effect of $\ln(R\&D_{i,t})$ is no longer comparable to those from Panels A-C: the point estimate on total R&D expenditure per capita (column 1) is both economically small and statistically insignificant while the estimate on Business (column 3) switches sign all together. That on H. Edu. and Govt. remain positive and significant (columns 2 and 4), however, are smaller in magnitude. Based on this, we conclude that, consistent with the use of PPML for standard gravity regressions, it is probably preferable to also rely on it for second-stage analysis of the type that we describe here.

5 Recovering Trade Elasticities

The trade elasticity is a crucial parameter in international economics. In the trade literature, not only is it of paramount importance for calibration of GE models, but its associated value has large implications for welfare estimates (Arkolakis et al., 2012; Feenstra et al., 2018). In the macro literature, the trade elasticity parameter is also key because it determines the

sign and size of international spillovers (Lisak et al., 2021). In both literatures, the trade elasticity is thus non-trivial to drawing accurate conclusions and informing optimal policy recommendations.

Typically, however, researchers select a given elasticity (or elasticities) from the literature for model calibration, rather than estimating it themselves. And there are a plethora of estimates from which to choose. Because of its central role for international economic analysis, many studies have offered a range of methods to compute the trade elasticity, and provide elasticity estimates at various levels of aggregation. While surely a non-exhaustive list, some influential studies that examine this parameter include Broda and Weinstein (2006), Egger et al. (2012), Hillberry and Hummels (2013), Simonovska and Waugh (2014), Soderbery (2015), Caliendo and Parro (2015), and Fontagné et al. (2020); see Feenstra et al. (2018) for a review and comparison of approaches and elasticity magnitudes.

However, simply drawing upon an elasticity from the literature comes with two potential drawbacks. On the one hand, researchers' chosen elasticity might have been estimated using data sources which do not necessarily fully match the countries, periods, and/or sectors used in their models. On the other, and most importantly, a given elasticity often does not correspond to the level of (dis-)aggregation required. When the case, this could cause researchers to impose perhaps unrealistic assumptions about which elasticity should apply to more or less aggregated goods sectors, or to services sectors for which estimates are rarely readily available.

In this section, we show how the structure of our estimating framework can be applied to recover trade elasticities at both aggregate and sectoral levels. Overall, we view our approach to recover trade elasticities as complementary to the numerous studies which estimate the trade elasticity for two reasons: our methodology is both straightforward to implement and allows us to compute this key parameter without the need for tariff and/or price data. Taking our estimations to the sectoral level, we present elasticity estimates for Manufacturing and Services. We also split aggregate data into Tradable versus Non-tradable sectors. However, the generalizability of our methods implies that elasticities are possible to obtain at various levels of (dis-)aggregation. Recall from Section 2.1 that the trade elasticity, $1 - \sigma$, is directly linked to the elasticity of substitution, σ . For convenience, in what follows our discussion focuses on σ .

In principle, as discussed when describing the structural parameters of equation (15), one can directly recover the elasticity of substitution from the estimated parameter on the OMR, i.e. $\alpha_6 = (1 - \sigma)/\sigma$. When doing this for aggregate trade based on results from our baseline specification in column (6) of Table 1, we obtain an elasticity of substitution $\sigma = 1/(1 - 0.878) = 8.20$. However, this specification does not fully exploit the structure

of the parameters imposed by our model, which, guided by theory, might prove useful in better identifying the elasticity. As such, in what follows we exploit the full structure of the parameters when estimating our baseline specification. More specifically, as discussed in Section 2.2, the structural interpretation of the parameters under constant returns to scale (CRS) is as follows: $\alpha_6 = (1 - \sigma)/\sigma$; $\alpha_1 = -\alpha_6$; $\alpha_2 = -(1 - \gamma)\alpha_6$; and $\alpha_3 = -\gamma\alpha_6$, where γ is the output elasticity of capital.²⁷ However, since our model also holds in the absence of the CRS assumption, we estimate a more flexible version as well by defining $\alpha_2 = -\beta\alpha_6$, where β is the output elasticity of labor (and keeping all other constraints the same). When relaxing the CRS assumption, we also provide a test of whether it holds, i.e. we test the null hypothesis that $\beta + \gamma = 1$.

With this in mind, Table 5 presents recovered elasticities of substitution for aggregate trade. Columns (1) and (2) show elasticities for our baseline model under CRS and non-CRS, respectively (where $\tilde{\alpha}$ corresponds to α_6). A few observations are in order. First, the elasticity of substitution reported in column (2) is quite high relative to that in column (1), but with large standard errors, thus rendering any interpretation of this result challenging. Second, in column (2), we marginally fail to reject the null hypothesis of CRS. Third, compared to our baseline results in column (6) of Table 1, the point estimate on the OMR term, $\tilde{\alpha}$, is identical but with standard errors which are close to double in magnitude. This potentially points to the difficulties in identifying the parameters in this constrained version of the model. Finally, in both columns we see that the elasticity of capital, γ , is relatively small in magnitude and statistically insignificant.

Continuing, columns (3) and (4) present results when implementing the alternative procedure to control for productivity described both above and in Supplementary Appendix A. The advantage of this method is that, instead of using an ex-ante estimate of TFP as in columns (1) and (2), we can directly account for unobserved productivity within our estimation. Hence, implementing this alternative procedure controls for TFP without the need for productivity data. As can be observed, irrespective of whether or not the CRS assumption is imposed, we recover very similar estimates for the output elasticity of capital, γ , and $\tilde{\alpha}$. The latter result implies an aggregate elasticity of substitution, σ , of 6.5, which is a bit larger than some recent trade elasticity estimates, e.g., Simonovska and Waugh (2014) and Caliendo and Parro (2015), but readily comparable to the majority of estimates from the trade literature, e.g., Anderson and van Wincoop (2003), Egger et al. (2012), Head and Mayer (2014), and Fontagné et al. (2020).

Interestingly, we fail to reject the null hypothesis of CRS in column (4), which lends support to the robustness of this approach. When comparing with column (6) of Table

²⁷By definition, $1 - \gamma$ is the output elasticity of labor.

3—whereby we also implement the alternative procedure but do not impose constraints on the parameters of our baseline specification—we note that the point estimate on the OMR term is strikingly similar. Finally, the estimated output elasticities of capital and labor are within the expected range. We therefore conclude that the alternative estimation approach (i) delivers robust results, and (ii) controls for unobserved productivity in a consistent manner. As we shall see below, this is particularly important when taking results to the sectoral level due to paucity of disaggregate TFP estimates.

Table 5: Aggregate Trade Elasticities

	(1)	(2)	(3)	(4)
	Baseline		Alternative Procedure	
	CRS	no CRS	CRS	no CRS
$\tilde{\alpha}$	-0.719*** (0.025)	-0.878*** (0.099)	-0.845*** (0.014)	-0.846*** (0.024)
γ	0.119 (0.101)	0.101 (0.085)	0.327*** (0.024)	0.322*** (0.029)
β		0.756*** (0.121)		0.677*** (0.024)
σ	3.555*** (0.312)	8.218 (6.654)	6.450*** (0.578)	6.509*** (1.004)
$\gamma + \beta$		0.857*** (0.069)		0.999*** (0.015)
$H_0 : \beta + \gamma = 1$ (p-value)	n/a	0.038	n/a	0.950
Observations	26,460	26,460	24,696	24,696

Notes: $\tilde{\alpha} = (1 - \sigma)/\sigma$ where σ is the elasticity of substitution. γ is the output elasticity of capital, β is the output elasticity of labor which is only estimated in the models which do not impose constant return to scale (CRS). “n/a” stands for not applicable. Standard errors clustered by country-pair in parentheses. * $p < 0.10$, ** $p < .05$, *** $p < .01$. See text and Supplementary Appendix A for further details.

Next, in Table 6 we bring our analysis to the sectoral level using the preferred alternative estimating procedure. Specifically, we present results for: Manufacturing and Services; as well as Tradable and Non-tradable sectors.²⁸ Panel A presents results when imposing CRS

²⁸See Supplementary Appendix Table B1 for a detailed breakdown of all WIOD industries. The Manufacturing sector encompasses WIOD industries 5-23 while the Services sector encompasses WIOD industries 28-50. Following Piton (2021), the Tradable sector includes (WIOD industries in parentheses): Mining and quarrying (4); Manufacturing (5-23); Accommodation and food service activities (36); Transportation and storage (31-35); Professional, scientific and technical activities (45-49); Administrative and support service activities (45); Information and communication (37-40); and Financial and insurance activities (41-43). The Non-tradable sector includes: Construction (27); Utilities and waste management (24-26); and Wholesale and retail trade, and repair of vehicles (28-30). Piton (2021) also excludes from her definition of Tradable and Non-tradable: all sectors which are part of the non-market economy (51-56); real estate activities (44) due to the fact that this sector is mostly composed of rental income; and agriculture (1-3), for which revenue is driven by European subsidies.

while Panel B relaxes this assumption. For ease of comparison, in column (1) we reproduce aggregate results from the last two columns of Table 5.

When comparing point estimates of $\tilde{\alpha}$ and γ in Panels A and B, we find comparable estimates except for column (5). This latter difference seems to be driven by the fact that the CRS assumption imposed in Panel A fails to hold for the specification in column (5), as suggested when testing the null hypothesis (H_0) of CRS at the bottom of Panel B. We therefore rely upon Panel B when interpreting our estimates. We find for manufacturing an elasticity of substitution $\hat{\sigma}_{Manufacturing} = 5.28$ (column 2), which is in line with the average estimate of 5.5 for goods from Fontagné et al. (2020).

Table 6: Sectoral Trade Elasticities

	(1)	(2)	(3)	(4)	(5)
	Aggregate	Manufacturing	Services	Goods & Services Tradable	Services Sectors: Non-tradable
Panel A: CRS					
$\tilde{\alpha}$	-0.845*** (0.014)	-0.851*** (0.017)	-0.898*** (0.016)	-0.819*** (0.016)	-0.844*** (0.019)
γ	0.327*** (0.024)	0.440*** (0.016)	0.362*** (0.025)	0.444*** (0.017)	0.184*** (0.032)
σ	6.450*** (0.578)	6.719*** (0.789)	9.797*** (1.565)	5.537*** (0.501)	6.408*** (0.772)
Panel B: no CRS					
$\tilde{\alpha}$	-0.846*** (0.024)	-0.810*** (0.047)	-0.868*** (0.027)	-0.810*** (0.032)	-0.684*** (0.034)
γ	0.322*** (0.029)	0.458*** (0.027)	0.408*** (0.042)	0.451*** (0.031)	0.383*** (0.057)
β	0.677*** (0.024)	0.568*** (0.018)	0.630*** (0.027)	0.558*** (0.018)	0.853*** (0.038)
σ	6.509*** (1.004)	5.276*** (1.308)	7.602*** (1.567)	5.254*** (0.884)	3.164*** (0.344)
$\gamma + \beta$	0.999*** (0.015)	1.026*** (0.029)	1.038*** (0.028)	1.009*** (0.024)	1.237*** (0.051)
$H_0 : \beta + \gamma = 1$ (p-value)	0.950	0.374	0.171	0.710	0.000
Observations	24,696	24,696	24,696	24,696	24,696

Notes: $\tilde{\alpha} = (1 - \sigma)/\sigma$ where σ is the elasticity of substitution, γ is the output elasticity of capital, β is the output elasticity of labor which is only estimated in the models which do not impose constant return to scale (CRS). Standard errors clustered by country-pair in parentheses. * $p < 0.10$, ** $p < .05$, *** $p < .01$. For further estimation details, see Supplementary Appendix A.

Column (3) presents our results for Services. Overall, we find an elasticity of substitution $\hat{\sigma}_{Services} = 7.60$, which is roughly 45 percent larger than $\hat{\sigma}_{Manufacturing}$. We find this result intuitive given the nature of services trade. Importantly, this estimate is higher than what previous literature which has attempted to estimate the services elasticity suggests. For

instance, Egger et al. (2012) obtain an elasticity estimate $\hat{\sigma}_{Services} = 5.96$ after developing a structural model of goods and services trade and production and implementing a systems estimation strategy to identify all parameters. Thus, our novel results hint that this sector might be even more substitutable than we previously thought.

Next, in columns (4) and (5) we examine one additional slice of the data by estimating elasticities of substitution for Tradable versus Non-tradable sectors. This is an important exercise given that (i) some economic sectors exhibit minimal trade intensities, and (ii) this exercise can be used to further validate our results since the elasticities for Tradables are expected to be larger. In line with our priors, we find a Tradable elasticity $\hat{\sigma}_{Tradable} = 5.25$, which is more or less in line with that for manufacturing and roughly 65 percent larger than that for Non-tradables.

In sum, results in this section demonstrate the validity of our proposed methods to identify the elasticity of substitution (σ), and hence the trade elasticity ($1 - \sigma$), at the aggregate and sectoral level within a structural gravity setting. One large merit of this approach is that, provided availability of the necessary production and trade data, researchers can directly implement our methods to estimate their own elasticities at their preferred levels of (dis-)aggregation and for their relevant country and time samples without the need for tariff and/or price data.

6 Concluding Remarks and Avenues for Future Work

Complementing theory with empirics, we derive a simple two-stage estimating procedure which is straightforward to estimate and which contributes to the empirical trade literature across two important dimensions.

First, we contribute to the gravity literature, which aims to quantify the effects of various determinants of trade flows, by proposing methods to identify the full impact of country-specific characteristics on international trade. While various methods to date have been proposed to do this, our methodology differs in that it uses bilateral trade flows (the key outcome of interest) as the second-stage dependent variable, and delivers estimates of country-specific determinants in a theoretically grounded way.

As proof of concept, we quantify the impact of research and development (R&D) expenditure on bilateral trade flows. Overall, we show that there is a strong link between gross R&D expenditure and international and domestic trade; *ceteris paribus*, a 10 percent increase in total R&D expenditure is associated with a 3 percent increase in bilateral trade flows. However, this result varies by type of R&D expenditure. Interestingly, government allocations of R&D spending promote domestic trade, but lead to lower international trade.

Importantly, while we demonstrate the viability of our methods through an examination of the impacts of R&D expenditure on bilateral trade flows, our approach can be applied to any country-specific variable, thus allowing economists to estimate the full effect of their country-specific determinant of choice in a theory-consistent manner.

Second, we contribute to the literature that empirically estimates trade elasticities, in our setting given by $1 - \sigma$, which is directly linked to the the elasticity of substitution, σ . In this regard, the most novel element of our methodology is that, through our estimating equation, we are able to recover the elasticity without the need for price and/or tariff data. Taking our strategy to the sectoral level, an important implication of our analysis is that we compute the elasticity of substitution for Services ($\hat{\sigma}_{Services} = 7.60$) which is roughly 45 percent larger than that for Manufacturing. Further, we demonstrate that while the elasticity of substitution for Tradable sectors ($\hat{\sigma}_{Tradable} = 5.25$) is roughly in line with that for Manufacturing, that for Non-tradable sectors ($\hat{\sigma}_{Non-tradable} = 3.16$) is, unsurprisingly, much smaller.

More generally speaking, our estimation strategy allows researchers to recover one of the most crucial parameters in the international economics literature, the trade elasticity, through a simple structural approach. Given that the gravity framework is frequently used for general equilibrium (GE) analysis, a powerful implication of our methods is that they can enable researchers to recover their own trade elasticities while using the same data applied in their counterfactual exercises and which correspond to the exact level of (dis-)aggregation required.

In sum, while our theoretical insights are, in many respects, the natural progression of four generations of gravity literature, the empirical possibilities they unlock are large. Looking to the future, we see many avenues for fruitful extensions to this work. In particular, decomposing the structural inward multilateral resistance term would allow to isolate the impact of international versus domestic links, and incorporate a structural estimating framework that simultaneously accounts for sectors and intermediates. This would naturally open the door to further apply our methods to the fast-moving literatures which address global value chains, structural gravity, and policy-induced GE/welfare effects.

References

- Akcigit, U., S. T. Ates, and G. Impullitti (2018). Innovation and trade policy in a globalized world. *NBER Working Papers No. 24543*.
- Anderson, J. E. (1979). A theoretical foundation for the gravity equation. *American Economic Review* 69(1), 106–116.
- Anderson, J. E. (2011). The gravity model. *Annual Review of Economics* 3, 133–160.
- Anderson, J. E., M. Larch, and Y. V. Yotov (2018). GEPPML: General equilibrium analysis with PPML. *The World Economy* 41(10), 2750–2782.
- Anderson, J. E., M. Larch, and Y. V. Yotov (2020). Transitional growth and trade with frictions: A structural estimation framework. *Economic Journal* 130(630), 1583–1607.
- Anderson, J. E. and E. van Wincoop (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic Review* 93(1), 170–192.
- Anderson, J. E. and E. van Wincoop (2004). Trade costs. *Journal of Economic Literature* 42(3), 691–751.
- Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012). New trade models, same old gains? *American Economic Review* 102(1), 94–130.
- Armington, P. (1969). A theory of demand for products distinguished by place of production. *IMF Staff Papers* 16, 159–176.
- Arvis, J.-F. and B. Shepherd (2013). The Poisson quasi-maximum likelihood estimator: A solution to the “adding up” problem in gravity models. *Applied Economics Letters* 20(6), 515–519.
- Aw, B. Y., M. J. Roberts, and D. Y. Xu (2011). R&D investment, exporting, and productivity dynamics. *American Economic Review* 101(4), 1312–1344.
- Baier, S. L., A. Kerr, and Y. V. Yotov (2018). Gravity, distance, and international trade. In B. A. Blonigen and W. W. Wilson (Eds.), *Handbook of International Trade and Transportation*, Chapters, Chapter 2, pp. 15–78. Edward Elgar Publishing.
- Baldwin, R. and D. Taglioni (2006). Gravity for dummies and dummies for gravity equations. *NBER Working Paper No. 12516*.
- Beverelli, C., A. Keck, M. Larch, and Y. V. Yotov (2018). Institutional quality, trade and development: A quantitative analysis. *CESifo Working Paper No. 6920*.
- Blundell, R. and S. Bond (2000). GMM estimation with persistent panel data: An application to production functions. *Econometric Reviews* 19(3), 321–340.
- Broda, C. and D. Weinstein (2006). Globalization and the gains from variety. *Quarterly Journal of Economics* 121(2), 541–585.

- Caliendo, L. and F. Parro (2015). Estimates of the trade and welfare effects of NAFTA. *The Review of Economic Studies* 82(1), 1–44.
- Costinot, A., D. Donaldson, and I. Komunjer (2012). What goods do countries trade? A quantitative exploration of Ricardo’s ideas. *Review of Economic Studies* 79(2), 581–608.
- Costinot, A. and A. Rodriguez-Clare (2014). Trade theory with numbers: Quantifying the consequences of globalization. In G. Gopinath, E. Helpman, and K. S. Rogoff (Eds.), *Handbook of International Economics*, Volume 4, Chapter 4, pp. 197–261. Elsevier.
- Eaton, J. and S. Kortum (2002). Technology, geography and trade. *Econometrica* 70(5), 1741–1779.
- Eaton, J., S. Kortum, B. Neiman, and J. Romalis (2016). Trade and the global recession. *American Economic Review* 106(11), 3401–38.
- Egger, P. and M. Larch (2008). Interdependent preferential trade agreement memberships: An empirical analysis. *Journal of International Economics* 76(2), 384–399.
- Egger, P., M. Larch, and K. Staub (2012). Trade preferences and bilateral trade in goods and services: A structural approach. *CEPR Working Paper No. 9051*.
- Fally, T. (2015). Structural gravity and fixed effects. *Journal of International Economics* 97(1), 76–85.
- Feenstra, R. C., P. Luck, M. Obstfeld, and K. N. Russ (2018). In search of the Armington elasticity. *Review of Economics and Statistics* 100(1), 135–150.
- Feyrer, J. (2019). Trade and income—exploiting time series in geography. *American Economic Journal: Applied Economics* 11(4), 1–35.
- Feyrer, J. (2021). Distance, trade, and income—the 1967 to 1975 closing of the Suez canal as a natural experiment. *Journal of Development Economics* 153, 102708.
- Fontagné, L. G., H. Guimbard, and G. Orefice (2020). Product-level trade elasticities: Worth weighting for. *CESifo Working Paper Series No. 8491*.
- Grossman, G. M., P. McCalman, and R. W. Staiger (2021). The “new” economics of trade agreements: From trade liberalization to regulatory convergence? *Econometrica* 89(1), 215–249.
- Haaland, J. I. and H. J. Kind (2008). R&D policies, trade and process innovation. *Journal of International Economics* 74(1), 170–187.
- Head, K. and T. Mayer (2014). Gravity equations: Workhorse, toolkit, and cookbook. In G. Gopinath, E. Helpman, and K. S. Rogoff (Eds.), *Handbook of International Economics*, Volume 4, Chapter 3, pp. 131–195. Elsevier.

- Heid, B., M. Larch, and Y. V. Yotov (2021). Estimating the effects of non-discriminatory trade policies within structural gravity models. *Canadian Journal of Economics* 54(1), 376–409.
- Hillberry, R. and D. Hummels (2013). Trade elasticity parameters for a computable general equilibrium model. In P. B. Dixon and D. Jorgenson (Eds.), *Handbook of Computable General Equilibrium Modeling*, Volume 1, Chapter 18, pp. 1213–1269. Elsevier.
- Hummels, D. (2001). Toward a geography of trade costs. unpublished manuscript.
- Kasterine, A. (2015). The new world trade order is about “precaution not protection” – Pascal Lamy. Technical report.
- Lamy, P. (2012). The universe of non-tariff measures has changed. Technical report.
- Lisak, N., S. Lloyd, and R. Sajedi (2021). Aggregation across each nation: Aggregator choice and macroeconomic dynamics. *Mimeo*.
- Olivero, M. and Y. V. Yotov (2012). Dynamic gravity: Endogenous country size and asset accumulation. *Canadian Journal of Economics* 45(1), 64–92.
- Olley, G. and A. Pakes (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica* 64(6), 1263–1297.
- Peters, B., M. J. Roberts, and V. A. Vuong (2018). Firm R&D investment and export market exposure. *NBER Working Papers No. 25228*.
- Piton, S. (2021). Economic integration and unit labour costs. *European Economic Review* 136, 103746.
- Redding, S. and A. Venables (2004). Economic geography and international inequality. *Journal of International Economic* 62(1), 53–82.
- Santos Silva, J. M. C. and S. Tenreyro (2006). The log of gravity. *Review of Economics and Statistics* 88(4), 641–658.
- Santos Silva, J. M. C. and S. Tenreyro (2011). Further simulation evidence on the performance of the Poisson pseudo-maximum likelihood estimator. *Economics Letters* 112(2), 220–222.
- Simonovska, I. and M. E. Waugh (2014). The elasticity of trade: Estimates and evidence. *Journal of International Economics* 92(1), 34–50.
- Soderbery, A. (2015). Estimating import supply and demand elasticities: Analysis and implications. *Journal of International Economics* 96(1), 1–17.
- United Nations, S. D. (2008). International Standard Industrial Classification of All Economic Activities ISIC Rev. 4. https://unstats.un.org/unsd/publication/seriesm/seriesm_4rev4e.pdf. Statistical papers, Series M No. 4/Rev.4.

- Windmeijer, F. (2008). GMM for panel data count models. In *The Econometrics of Panel Data—Fundamentals and Recent Developments in Theory and Practice* (3 ed.), Chapter 18, pp. 603–624. Springer.
- Windmeijer, F. and J. Santos Silva (1997). Estimation of count data models with endogenous regressors: An application to demand for health care. *Journal of Applied Econometrics* 12(3), 281–294.
- Yotov, Y. V., R. Piermartini, J.-A. Monteiro, and M. Larch (2016). *An Advanced Guide to Trade Policy Analysis: The Structural Gravity Model*. Geneva: UNCTAD and WTO.

Supplementary Appendix

A Alternative Estimation Process for Productivity

This Appendix details our estimating methodology for the alternative approach to account for unobserved productivity described in the main text. We see this as a potentially important contribution for two reasons. First, the TFP term used in our baseline specification from PWT is based on estimates of a production function which does not necessarily align with all underlying assumptions in our structural framework and is estimated using different data sources than those in our main analysis. Second, we do not observe TFP at the sectoral level, which is problematic when estimating trade elasticities for Manufacturing, Services and Tradable versus Non-tradable sectors. Hence, the approach described below allows us to bypass this concern through alleviating the need to rely upon external estimates of TFP.

Step 1 remains identical to that detailed in the main text (i.e. equation (16) in Section 2.2) whereby we obtain estimates for: the power transformation of the trade cost vector $\left(\widehat{t_{ij,t}^{1-\sigma}}\right)$; importer-time fixed effects $\left(\widehat{\psi_{j,t}}\right)$; exporter-time fixed effects $\left(\widehat{\phi_{i,t}}\right)$, which allow us to compute the power transformation of the OMR term $\left(\widehat{\Pi_{i,t}^{1-\sigma}}\right)$; and the error term $\left(\widehat{\epsilon_{ij,t}}\right)$.

Subsequently, using our preferred baseline specification (as per column (6) of Table 1), we now net out the estimated (from step one) trade costs, importer-time fixed effects, and error term vectors from bilateral trade:

$$\widetilde{X}_{ij,t} = \exp \left[\alpha_1 \ln(A_{i,t}) + \alpha_2 \ln(L_{i,t}) + \alpha_3 \ln(K_{i,t}) + \alpha_7 \ln \left(\widehat{\Pi_{i,t}^{1-\sigma}} \right) + \phi_{i,t} \right], \quad (21)$$

where $\widetilde{X}_{ij,t} \equiv X_{ij,t} / \left(\widehat{t_{ij,t}^{1-\sigma}} \widehat{\psi_{j,t}} \widehat{\epsilon_{ij,t}} \right)$ and $\phi_{i,t}$ represent exporter and time fixed effects.

Since productivity is unobserved, we follow the dynamic panel and proxy variable methods from the productivity literature (Blundell and Bond, 2000; Olley and Pakes, 1996) in assuming that the power transformation of productivity $\ln \left(\widetilde{A}_{i,t} \right) = \alpha_1 \ln(A_{i,t})$ follows a Markov process, i.e. first order auto-regressive process, $\ln \left(\widetilde{A}_{i,t} \right) = \rho \ln \left(\widetilde{A}_{i,t-1} \right) + \xi_{it}$, where ξ_{it} reflects productivity shocks.

Under this assumption about the evolution of productivity, we can re-write equation (21) as follows:

$$\widetilde{X}_{ij,t} = \exp \left[\rho \ln \left(\widetilde{A}_{i,t-1} \right) + \alpha_2 \ln(L_{i,t}) + \alpha_3 \ln(K_{i,t}) + \alpha_7 \ln \left(\widehat{\Pi_{i,t}^{1-\sigma}} \right) + \phi_{i,t} \right] \times \widetilde{\xi}_{ij,t}, \quad (22)$$

where $\widetilde{\xi}_{ij,t} \equiv \exp[\xi_{i,t}]$. Then, taking the log of equation (22), we can re-express productivity

as:

$$\ln(\tilde{A}_{i,t}) = \ln(\tilde{X}_{ij,t}) - \alpha_2 \ln(L_{i,t}) - \alpha_3 \ln(K_{i,t}) - \alpha_7 \ln(\widehat{\Pi_{i,t}^{1-\sigma}}) - \phi_{i,t}. \quad (23)$$

This allows us to now substitute the lagged version of equation (23) into (22) to obtain:

$$\begin{aligned} X_{ij,t} = & \exp \left[\rho \left(\ln(\tilde{X}_{ij,t-1}) - \alpha_2 \ln(L_{i,t-1}) - \alpha_3 \ln(K_{i,t-1}) - \alpha_7 \ln(\widehat{\Pi_{i,t-1}^{1-\sigma}}) \right) \right] \\ & \times \exp \left[\alpha_2 \ln(L_{i,t}) + \alpha_3 \ln(K_{i,t}) + \alpha_7 \ln(\widehat{\Pi_{i,t}^{1-\sigma}}) + \tilde{\phi}_{i,t} \right] \times \tilde{\xi}_{ij,t}, \end{aligned} \quad (24)$$

We can now bring this equation to the data using Generalized Method of Moments (GMM), where moment conditions are based on the same variables used in equation (24) since we assume that the error term is independently distributed.²⁹ In addition, we use exporter and time dummies to control for the fixed effects which are scaled up by a constant, i.e. $\tilde{\phi}_{i,t} \equiv (1 - \rho)\phi_{i,t}$.³⁰ Finally, we can also estimate a version of equation (24) where we directly impose the parameter constraints implied by our model, i.e. $\alpha_3 = -\gamma\alpha_7$, where γ is the output elasticity of capital. Then, under constant returns to scale, $\alpha_2 = (\gamma - 1)\alpha_7$. Without constant returns to scale $\alpha_2 = -\beta\alpha_7$, where β is the output elasticity of labor.

²⁹Note that we can also allow for correlated shocks over time and across countries, but this would require the use of deeper lags for instruments given this dynamic setting. This would pose a limitation in our current setup given the relatively short time dimension of the WIOD dataset. Further, recall that we have already accounted for correlated shocks in the first stage error term, $\epsilon_{ij,t}$, which helps to alleviate concerns of this nature.

³⁰See Windmeijer (2008) for an overview of the literature on estimating dynamic models for panel count data.

B WIOD Data Description

Table B1: WIOD Sector Descriptions and Broad Sector Groupings

ISIC R4 Code	WIOD ID	Detailed Description	Broad Sector Grouping
A	1-3	Agriculture	Agriculture
B	4	Mining & quarrying	Mining & Quarrying
C	5-23	Manufacturing	Manufacturing
D & E	24-26	Utilities & Waste Management	Utilities & Waste
F	27	Construction	Construction
G	28-30	Wholesale & retail trade; repair of vehicles	Services
H	31-35	Transportation & storage	Services
I	36	Accommodation & food service activities	Services
J	37-40	Information & communication	Services
K	41-43	Financial & insurance activities	Services
L	44	Real estate activities	Services
M	45-49	Professional, scientific & technical activities	Services
N	50	Administrative & support service activities	Services
O	51	Public admin. & defence; compulsory social security	Non-market Economy
P	52	Education	Non-market Economy
Q	53	Human health & social work activities	Non-market Economy
R & S	54	Arts, entertainment & recreation; Other services	Non-market Economy
T	55	Activities of households as employers	Non-market Economy
U	56	Activities of extraterritorial organizations & bodies	Non-market Economy

Notes: Broad sector groupings are classified as per the broad sector concordance with ISIC Rev.4 (United Nations, 2008).

Table B2: Countries Included in Analysis

ISO-3 Code	Country Name
AUS	Australia
AUT	Austria
BEL	Belgium
BGR	Bulgaria
BRA	Brazil
CAN	Canada
CHE	Switzerland
CHN	China
CYP	Cyprus
CZE	Czech Republic
DEU	Germany
DNK	Denmark
ESP	Spain
EST	Estonia
FIN	Finland
FRA	France
GBR	United Kingdom
GRC	Greece
HRV	Croatia
HUN	Hungary
IDN	Indonesia
IND	India
IRL	Ireland
ITA	Italy
JPN	Japan
KOR	Korea
LTU	Lithuania
LUX	Luxembourg
LVA	Latvia
MEX	Mexico
MLT	Malta
NLD	Netherlands
NOR	Norway
POL	Poland
PRT	Portugal
ROU	Romania
RUS	Russia
SVK	Slovak Republic
SVN	Slovenia
SWE	Sweden
TUR	Turkey
USA	United States

C Full Country-specific Results

Table C1: Full Country-specific Results Using Two-stage Structural Gravity Estimation Procedure

	(1)	(2)	(3)	(4)
	Total	H. Edu.	Business	Govt.
Stage 1				
$RTA_{ij,t}$	0.071 (0.058)	0.100* (0.057)	0.046 (0.059)	0.045 (0.049)
$EU_{ij,t}$	0.068** (0.033)	0.038 (0.031)	0.066** (0.033)	0.148*** (0.037)
$WTO_{ij,t}$	0.040 (0.050)	0.038 (0.050)	0.040 (0.050)	-0.157*** (0.051)
$\ln(R\&D_{i,t}) * BRDR_{ij}$	0.244*** (0.094)	0.115* (0.059)	0.113* (0.060)	-0.176*** (0.049)
Observations	20,958	21,168	21,126	20,202
Stage 2				
$\ln(A_{i,t})$	0.039 (0.029)	0.019 (0.043)	0.049 (0.034)	-0.006 (0.034)
$\ln(L_{i,t})$	0.495*** (0.048)	0.526*** (0.039)	0.504*** (0.053)	0.564*** (0.043)
$\ln(K_{i,t})$	0.396*** (0.040)	0.403*** (0.045)	0.402*** (0.044)	0.366*** (0.037)
$\ln(\widehat{\Pi_{i,t}^{1-\sigma}})$	-0.860*** (0.057)	-0.886*** (0.052)	-0.846*** (0.064)	-0.879*** (0.059)
$\ln(R\&D_{i,t})$	0.081*** (0.014)	0.075*** (0.019)	0.051*** (0.013)	0.082*** (0.018)
Total effect	0.325	0.191	0.164	-0.094
Observations	20,958	21,168	21,126	20,202

Notes: The country-specific variable we include is gross domestic expenditure on R&D at current PPPs per capita (column 1), broken down into sub-categories for: higher education (H. Edu., column 2); business enterprise expenditure (Business, column 3); and total government allocations (Govt., column 4). Standard errors clustered by country-pair in parentheses. * $p < 0.10$, ** $p < .05$, *** $p < .01$.

Table C2: Full Country-specific Results Under Naïve Gravity

	(1)	(2)	(3)	(4)
	Total	H. Edu.	Business	Govt.
$RTA_{ij,t}$	0.088* (0.048)	0.115** (0.047)	0.070 (0.056)	0.057 (0.043)
$EU_{ij,t}$	0.136*** (0.041)	0.116*** (0.042)	0.142*** (0.043)	0.281*** (0.041)
$WTO_{ij,t}$	0.022 (0.062)	0.022 (0.061)	0.023 (0.061)	-0.184*** (0.057)
$\ln(\text{remote}_{i,t})$	0.180 (0.160)	-0.151 (0.174)	-0.065 (0.184)	0.337** (0.167)
$\ln(GDP_{i,t})$	0.568*** (0.045)	0.693*** (0.055)	0.677*** (0.044)	0.646*** (0.049)
$\ln(R\&D_{i,t}) * BRDR_{ij}$	0.232** (0.095)	0.114** (0.058)	0.104 (0.072)	-0.139** (0.058)
$\ln(R\&D_{i,t})$	0.082 (0.055)	0.016 (0.032)	0.039 (0.041)	0.069* (0.038)
Total effect	0.314	0.131	0.143	-0.070
Observations	20,958	21,168	21,126	20,202

Notes: The country-specific variable we include is gross domestic expenditure on R&D at current PPPs per capita (column 1), broken down into sub-categories for: higher education (H. Edu., column 2); business enterprise expenditure (Business, column 3); and total government allocations (Govt., column 4). Standard errors clustered by country-pair in parentheses. * $p < 0.10$, ** $p < .05$, *** $p < .01$.

Table C3: Full Country-specific Stage-two Results Under Approach Proposed in Head and Mayer (2014)

	(1)	(2)	(3)	(4)
	Total	H. Edu.	Business	Govt.
FE PPML				
$\ln(A_{i,t})$	0.100*** (0.039)	0.073 (0.057)	0.104** (0.049)	0.045 (0.028)
$\ln(L_{i,t})$	0.567*** (0.059)	0.583*** (0.043)	0.575*** (0.060)	0.564*** (0.046)
$\ln(K_{i,t})$	0.344*** (0.052)	0.354*** (0.051)	0.354*** (0.055)	0.330*** (0.042)
$\ln(\widehat{\Pi_{i,t}^{1-\sigma}})$	-0.928*** (0.079)	-0.924*** (0.068)	-0.912*** (0.081)	-0.868*** (0.055)
$\ln(R\&D_{i,t})$	0.069*** (0.021)	0.046*** (0.017)	0.031* (0.019)	0.061*** (0.021)
Total effect	0.313	0.161	0.144	-0.115
Observations	499	504	503	481
FE OLS				
$\ln(A_{i,t})$	0.094 (0.068)	0.108** (0.050)	0.098 (0.069)	0.067 (0.060)
$\ln(L_{i,t})$	0.562*** (0.057)	0.558*** (0.051)	0.564*** (0.056)	0.540*** (0.060)
$\ln(K_{i,t})$	0.343*** (0.058)	0.332*** (0.056)	0.355*** (0.055)	0.357*** (0.054)
$\ln(\widehat{\Pi_{i,t}^{1-\sigma}})$	-0.887*** (0.062)	-0.905*** (0.057)	-0.886*** (0.058)	-0.899*** (0.063)
$\ln(R\&D_{i,t})$	0.014 (0.025)	0.037** (0.014)	-0.003 (0.014)	0.059** (0.027)
Total effect	0.258	0.152	0.110	-0.117
Observations	499	504	503	481

Notes: The country-specific variable we include is gross domestic expenditure on R&D at current PPPs per capita (column 1), broken down into sub-categories for: higher education (H. Edu., column 2); business enterprise expenditure (Business, column 3); and total government allocations (Govt., column 4). First-stage results are presented in Table C1. Standard errors clustered by country-pair in parentheses. * $p < 0.10$, ** $p < .05$, *** $p < .01$.