THE
ECONOMIC
JOURNAL

The Economic Journal, **132** (*May*), 1349–1377 https://doi.org/10.1093/ej/ueab067
Advance Access Publication Date: 4 September 2021

GRAVITY AND HETEROGENEOUS TRADE COST ELASTICITIES*

Natalie Chen and Dennis Novy

How do trade costs affect international trade? This paper offers a new approach. We rely on a flexible gravity equation that predicts variable trade cost elasticities, both across and within country pairs. We apply this framework to popular trade cost variables such as currency unions, trade agreements and World Trade Organization membership. While we estimate that these variables are associated with increased bilateral trade on average, we find substantial heterogeneity. Consistent with the predictions of our framework, trade cost effects are strong for 'thin' bilateral relationships characterised by small import shares, and weak or even zero for 'thick' relationships.

A key research topic in international trade is to understand the link between trade costs and trade flows. In this paper, we propose a new approach that is built on the idea that trade costs may not affect all trade flows in the same way. Instead, trade costs might have a strong influence on trade between some countries but not between others.

To evaluate the effect of trade costs, researchers typically rely on a standard gravity equation framework and insert trade cost proxies as right-hand side regressors (e.g., bilateral distance, dummy variables for regional trade agreements, currency union status, etc.) This yields single coefficients to assess the trade effects of these variables. By construction, their effects are homogeneous across all country pairs in the sample.[1]

In this paper, we challenge the view that trade costs have a homogeneous 'one-size-fits-all' effect on bilateral trade flows. The core of our paper is to offer an easy-to-implement alternative

* Corresponding author: Natalie Chen, Department of Economics, University of Warwick, Coventry CV4 7AL, UK. Email: n.a.chen@warwick.ac.uk

This paper was received on 14 June 2019 and accepted on 26 August 2021. The Editor was Nezih Guner.

The data and codes for this paper are available on the Journal repository. The data and codes for this paper were checked for their ability to reproduce the results presented in the paper. The authors were granted an exemption to publish parts of their data because access to these data is restricted. However, the authors provided the Journal with temporary access to the data, which enabled the Journal to run their codes. The codes for the parts subject to exemption are also available on the Journal repository. The restricted access data and these codes were also checked for their ability to reproduce the results presented in the paper. Given the highly demanding nature of the algorithms, the reproducibility checks were run on a simplified version of the code. The replication package for this paper is available at the following address: https://doi.org/10.5281/zenodo.5195761.

This paper was previously circulated as 'Currency Unions, Trade, and Heterogeneity.' We thank the editor Nezih Guner and two anonymous referees for many valuable comments. We gratefully acknowledge financial support from the Centre for Competitive Advantage in the Global Economy at the University of Warwick (CAGE, ESRC grant ES/L011719/1) and the Economic and Social Research Council (ESRC grant ES/P00766X/1). We are also grateful to Daniel Baumgarten, Douglas Campbell, Swati Dhingra, James Fenske, Boris Georgiev, Chris Meissner, John Morrow, Toshihiro Okubo, Andrew Rose, João Santos Silva and to participants at the China IEFS Conference 2017, the 13th Danish International Economics Workshop Aarhus 2017, the Royal Economic Society Conference Bristol 2017, the Society for International Trade Theory Conference Dublin 2017, the Workshop on Public Economic Policy Responses to International Trade Consequences Munich 2017, the Trade and Development Conference Hong Kong 2016, the Dynamics, Economic Growth, and International Trade Conference Nottingham 2016, the RIETI Conference Tokyo 2016, the Delhi School of Economics, the European Bank for Reconstruction and Development, the Paris Trade Seminar, Kiel, LSE, Middlesex, Nottingham, Oslo, Sussex, and to the CAGE and macro groups at Warwick.

[1] To be precise, the direct (partial equilibrium) trade effects are homogeneous. We discuss general equilibrium effects in Online Appendix B.

to the traditional gravity equation that allows us to estimate heterogeneous trade cost effects. Our contribution is to empirically demonstrate in a systematic and comprehensive way that trade cost effects are heterogeneous across country pairs, and also within country pairs by direction of trade.

As a prominent example, consider the trade effect of currency unions. Currency unions are arguably an important institutional arrangement to reduce trade costs. In the period since World War II, a total of 123 countries have been involved in a currency union at some point. By the year 2015, eighty-three countries continued to do so. In addition, various countries are currently considering to form new currency unions or to join existing ones.[2] But does that mean all currency union member pairs experience an equal increase in international trade? We provide an empirical framework showing that the trade effects associated with currency union membership are heterogeneous across member pairs, even within specific currency unions such as the euro.

We also apply this framework of flexible trade cost effects to a host of other trade cost-related variables popular in the international trade literature. We provide evidence of heterogeneous effects for regional trade agreements (RTAs), World Trade Organization (WTO) membership, bilateral distance, sharing a common border, a common language and tariffs.

As our theoretical framework, we introduce heterogeneous trade cost effects by taking guidance from a translog gravity equation that predicts variable trade cost elasticities (Novy, 2013). In this framework, 'thin' bilateral trade relationships (characterised by small bilateral import shares) are more sensitive to trade cost changes compared to 'thick' trade relationships (characterised by large bilateral import shares). The intuition is that small import shares are high up on the demand curve where sales are very sensitive to trade cost changes. Large import shares are further down on the demand curve where sales are more buffered. As a result, smaller import shares have a larger trade cost elasticity in absolute magnitude. The prediction is that a given change in trade costs (induced by a change in currency union status or other trade cost changes) generates heterogeneous effects on trade flows. We should expect larger trade effects for country pairs associated with smaller import shares. This implies a heterogeneous effect even within country pairs since bilateral import shares typically differ depending on the direction of trade.

We start by laying out the flexible gravity framework and relate it to trade cost effects in international trade. This forms the basis for our empirical specifications. We then construct our key variable of interest—the bilateral import shares of 199 countries between 1949 and 2013—and bring the framework to the data.[3] We adopt two main approaches to test whether the effect of trade costs on trade is heterogeneous across import shares. The first approach is a modification of the standard gravity specification familiar from the literature. Instead of estimating a single trade cost coefficient that is constant over the entire sample, we propose a flexible gravity framework that allows for heterogeneous trade cost effects across import shares.[4] The second approach is to estimate the translog gravity equation directly, which also implies variable trade cost elasticities.

In the first approach, our aim is to examine whether trade cost effects are heterogeneous across bilateral import shares. In principle, this could be attempted by interacting trade cost regressors

[2] Currency unions, or monetary unions, 'are groups of countries that share a single money' (Rose, 2006). See Online Appendix A for details. Areas currently considering the creation of a common currency include the economies of the West African Monetary Zone, the Southern African Development Community, the East African Community and the Gulf Cooperation Council (although in the latter case, talks have stalled).

[3] As explained in Section 1, the dependent variable is actually the bilateral import share per good of the exporting country. But, for simplicity, we refer to it as the bilateral 'import share'.

[4] It is well known that the gravity model fits the data very well. The main point of adopting the flexible gravity framework is not necessarily to improve overall fit but to introduce variable trade cost elasticities.

with import shares. However, this would create an immediate simultaneity bias problem since trade cost effects would vary with the values taken by the dependent variable. We address this issue by letting trade cost effects vary across *predicted* import shares. We propose a two-step methodology for this purpose. In the first step we generate predicted shares by regressing import shares on time-invariant geography-related variables such as distance and contiguity. In the second step we assess how the trade cost effects vary across predicted import shares. To deal with heteroskedasticity and to include the zero import shares in the sample, we estimate our regressions by Poisson Pseudo Maximum Likelihood (PPML), as is typical in the recent gravity literature.

We carry out Monte Carlo simulations to verify the validity of our two-step procedure of estimating heterogeneous trade cost effects. When we assume that the data generating process is driven by variable trade cost elasticities, we show that our two-step procedure produces heterogeneous results that match those implied by the model, both qualitatively and quantitatively. Also, when we assume that standard gravity with a constant trade cost elasticity is the data generating process, we demonstrate that our two-step procedure does not spuriously generate heterogeneous trade cost effects.

To illustrate our procedure, we initially focus on the effect of currency unions on international trade, and we turn to other trade cost variables later. Empirically, we find that, when we estimate a standard gravity regression without heterogeneous effects (i.e., imposing a constant trade cost elasticity), a common currency is associated with 29% more trade on average. Our contribution with the help of the flexible gravity framework is to demonstrate that this average hides a significant amount of heterogeneity *across country pairs*. Bilateral trade effects tend to be particularly strong between countries where at least one partner is relatively small so that import shares can be small. For example, we find a strong currency union effect for trade from Chad to Côte d'Ivoire (141% more bilateral trade) and from Austria to Germany (62%). Conversely, bilateral trade effects tend to be relatively weak or even zero between countries with large import shares, for example from Côte d'Ivoire to Togo (22%) and from France to Belgium-Luxembourg (insignificant).

We also find that the trade effect of currency unions is heterogeneous *within country pairs* and therefore asymmetric by direction of trade. For example, as just mentioned, the effect is large (62%) for trade from Austria to Germany (a low bilateral import share). But it is insignificant for trade from Germany to Austria (a large bilateral import share).

Given the enormous academic and policy interest in the euro, we also focus more specifically on the trade effect of the European single currency. Consistent with evidence reported in the literature, we confirm that the average trade effect of the euro is more modest compared to other common currencies.[5] This is consistent with relatively large import shares on average in the euro area. Still, we find that the effect is heterogeneous across country pairs within the Eurozone. Examples of country pairs with small import shares that are associated with large euro effects include Ireland importing from Cyprus (70%) and Finland importing from Malta (50%). Conversely, country pairs with large import shares not generating any additional trade from the common currency include Cyprus importing from Greece and Germany importing from Italy.

Recent contributions show that estimating a gravity equation by PPML as opposed to log-linear Ordinary Least Squares (OLS) reduces the size and significance of the estimated trade

[5] See Micco *et al.* (2003), Baldwin (2006), Baldwin and Taglioni (2007), Baldwin *et al.* (2008), Berger and Nitsch (2008), Santos Silva and Tenreyro (2010), Eicher and Henn (2011), Glick and Rose (2016), Mika and Zymek (2018), Larch *et al.* (2019), Mayer *et al.* (2019) or Campbell and Chentsov (2020).

effect of currency unions, and in particular of the euro.[6] Our framework contributes to explaining this finding. As is well known, OLS and PPML estimators have different first-order conditions (Eaton *et al.*, 2013; Head and Mayer, 2014; Mayer *et al.*, 2019). While the OLS conditions involve logarithmic deviations of trade from its expected value, the PPML conditions involve level deviations and therefore tend to give more weight to country pairs with high levels of trade. By highlighting that country pairs with higher trade intensity have a smaller trade cost elasticity, our framework therefore predicts that the PPML currency union estimate should be smaller than its OLS counterpart.

In the second approach, we explore the predictions of our model by estimating the translog gravity equation directly. Our regressions confirm that a common currency is associated with more bilateral trade, and that the magnitude of the effect falls with bilateral import shares.

One concern about our estimations relates to the potentially endogenous nature of trade cost variables. For example, reverse causality may arise because countries that trade intensively with each other are more likely to join a currency union, leading to an overestimation of the trade effect of common currencies.[7] Attempts in the literature to instrument the currency union dummy variable prove disappointing as the instrumentation tends to increase, rather than decrease, the magnitude of currency union estimates (Rose, 2000; Alesina *et al.*, 2002; Barro and Tenreyro, 2007). This has led the profession to conclude that appropriate instruments for currency union membership are not available (see Baldwin, 2006 for a discussion).

In this paper, we do not attempt to instrument the currency union indicator. But in simulation results we show that correcting for endogeneity bias (to the extent that it exists) should strengthen, rather than weaken, the heterogeneity patterns in our results. The intuition is that bilateral trade and currency unions are positively related. This would result in positive endogeneity bias, pushing up the modest currency union effects associated with large import shares. Thus, removing this potential bias would lead to even stronger heterogeneity patterns. Endogeneity can therefore not overturn our heterogeneity results.

As they improve our understanding of how trade costs shape trade flows between trading partners, our results have important policy implications. Most importantly, they help to evaluate the potential changes in bilateral trade flows that countries can expect when their trade costs change. For instance, suppose that Bulgaria, Croatia, the Czech Republic, Hungary, Poland, Romania and Sweden were to join the euro over the next few years. As these countries are relatively small compared to some members of the Eurozone such as France and Germany, they have relatively large import shares. Our results suggest that these import shares will grow only modestly (Baldwin, 2006; Glick, 2017; Mika and Zymek, 2018). However, trade shares in the opposite direction are smaller and can therefore be expected to grow more strongly.

Our approach is by no means the only one to explore heterogeneous trade cost elasticities. Novy (2013) concentrated on the theoretical derivation of the translog gravity framework, and he only explored its empirical implications based on bilateral distance and contiguity using a single cross section of twenty-eight OECD countries with fewer than a thousand observations.

[6] For instance, Santos Silva and Tenreyro (2010), Mika and Zymek (2018), Larch *et al.* (2019) and Mayer *et al.* (2019) found that PPML estimates of the euro trade effect are insignificant.

[7] Trading firms hurt by exchange rate fluctuations may lobby to keep the exchange rate with the country's major trading partners fixed (Baldwin, 2006). Reverse causality could also arise if currency unions capture unobserved characteristics that affect trade flows. For evidence that greater bilateral trade reduces bilateral exchange rate volatility, see Devereux and Lane (2003) and Broda and Romalis (2010). Mundell (1961) suggested that, by reducing real exchange rate fluctuations, trade reduces the costs of forming a currency union. Alesina and Barro (2002) showed that countries trading more with each other are more likely to form currency unions.

By contrast, we use more than 1.1 million observations and explore the empirical importance of heterogeneous trade cost effects in a more comprehensive way using a broad range of trade cost variables popular in the literature including time-varying components such as currency union status, regional trade agreements and WTO membership. Bas *et al.* (2017) derived country pair aggregate trade cost elasticities from a monopolistic competition model with Constant Elasticity of Substitution (CES) demand and log-normal firm-level heterogeneity. Consistent with our predictions, they found that the trade cost elasticity is smaller in magnitude for country pairs with large trade volumes. Guided by monopolistic competition models with CES demand and truncated Pareto firm-level productivity, Helpman *et al.* (2008) found that bilateral trade cost elasticities are larger for less developed countries, while Melitz and Redding (2015) documented that those elasticities vary across markets and levels of trade costs. Spearot (2013) relied on the Melitz and Ottaviano (2008) model with firm-level heterogeneity and linear demand to show that tariff liberalisation disproportionately increases the imports of low revenue varieties. Carrère *et al.* (2020) also stressed the importance of non-constant trade elasticities, focusing on the effect of distance in particular.[8]

More specifically, our paper also contributes to a large and growing literature that explores whether currency unions promote trade. In his seminal work, Rose (2000) showed that sharing a common currency more than triples bilateral trade flows. Subsequent work by Rose and co-authors showed that the currency union effect is smaller than initially found but remains large (Rose and van Wincoop, 2001; Frankel and Rose, 2002; Glick and Rose, 2002). Various authors argue that these findings are plagued by omitted variables, econometric errors and self-selection, may be driven by currency unions between small or poor countries and that the trade effect of currency unions is small or insignificant.[9] We rely on state-of-the-art PPML techniques that allow us to include zero trade observations in the sample and control for country pair and time-varying exporter and importer fixed effects. In addition, our results remain robust to endogeneity, self-selection, omitted variables such as wars and political conflicts, and to excluding small and poor countries as well as post-Soviet states from the sample.

Nevertheless, there is empirical evidence to suggest that heterogeneity in the trade impact of currency unions exists along several dimensions.[10] For instance, the effect is larger for developing economies (Santos Silva and Tenreyro, 2010), smaller countries (Micco *et al.*, 2003; Baldwin, 2006) and falls over time (De Sousa, 2012). The effect also varies across currency unions (Nitsch, 2002; Klein, 2005; Eicher and Henn, 2011; Glick and Rose, 2016). Consensus estimates for the euro tend to be more modest than those for broader samples, falling between 5% and 15% (Baldwin, 2006; Baldwin *et al.*, 2008). The trade effect is stronger for industries producing differentiated goods (Flam and Nordström, 2007), and for larger and more productive firms that adjust both at the intensive and extensive margins (Berthou and Fontagné, 2008). In contrast

[8] Also see Atkeson and Burstein (2008), who derived heterogeneous trade cost elasticities from a model with nested CES demand and oligopolistic competition.

[9] See Persson (2001), Nitsch (2002), De Nardis and Vicarelli (2003), López-Córdova and Meissner (2003), Micco *et al.* (2003), Klein (2005), Baldwin (2006), Klein and Shambaugh (2006), Bun and Klaassen (2007), Baldwin *et al.* (2008), Berger and Nitsch (2008), Broda and Romalis (2010), Frankel (2010), Santos Silva and Tenreyro (2010), Eicher and Henn (2011), De Sousa (2012), Campbell (2013), Glick and Rose (2016), Glick (2017), Saia (2017), Mika and Zymek (2018), Larch *et al.* (2019) and Campbell and Chentsov (2020). Baldwin *et al.* (2008) claimed that the empirical literature on the trade effect of currency unions 'is a disaster' as the estimates range from 0% (e.g., Berger and Nitsch, 2008) to 1,387% (Alesina *et al.*, 2002), most of them being 'fatally flawed by misspecification and/or econometric errors'.

[10] On the heterogeneous trade effects of FTAs, see Glick (2017) and Baier *et al.* (2019). Spearot (2013) studied the heterogeneous trade effects of tariff liberalisation. Subramanian and Wei (2007) and Felbermayr *et al.* (2020) explored the heterogeneous trade effects of WTO membership. Mayer *et al.* (2019) found that the trade effects of belonging to the EU are heterogeneous across member states, and the countries gaining the most are small, open and centrally located.

to these papers where the various sources of heterogeneity are explored without theoretical motivation and often across different samples, we are guided by a gravity framework with flexible trade cost elasticities to derive our empirical specifications.

The remainder of the paper is organised as follows. In Section 1 we build on the translog gravity framework to motivate why we might find heterogeneous trade cost effects in the data. In Section 2 we present our main estimation results. When introducing our empirical methodology, we initially focus on currency unions as a well-known trade cost variable to demonstrate the heterogeneity of trade effects. Then in Section 3 we extend the heterogeneity analysis to other prominent trade cost variables popular in the gravity literature such as RTAs, WTO membership, tariffs, bilateral distance and a shared border and language between trading partners. In Section 4 we carry out Monte Carlo simulations that explore the endogeneity of currency unions. In Section 5 we summarise an extensive battery of robustness checks. We conclude in Section 6. Online Appendix A summarises our data and sources. In Online Appendix B we outline the derivation of the translog gravity equation. We also carry out Monte Carlo simulations that validate our estimation strategy, and we discuss general equilibrium effects. Online Appendix C provides details on the robustness checks.

1. Theoretical Motivation

The conventional gravity framework in the literature is characterised by a constant trade cost elasticity. This means that the direct effect of a trade cost change is common across country pairs.[11] In this paper, we employ a gravity framework that does not feature a constant trade cost elasticity. Instead, we build on a gravity framework with variable trade cost elasticities. It follows that the effect of a trade cost change is no longer common across country pairs. It becomes heterogeneous.

As a framework that accommodates the crucial feature of variable trade cost elasticities, we use translog gravity to motivate our analysis. As in Novy (2013), the model features multiple countries in general equilibrium that are endowed with an arbitrary number of differentiated goods. Demand is derived from a translog expenditure function using the parameterisation in Feenstra (2003). Trade costs follow the iceberg form where $t_{ij} \geq 1$ denotes the bilateral trade cost factor between countries i and j. Trade costs may be bilaterally asymmetric such that $t_{ij} \neq t_{ji}$.

As outlined in Online Appendix B.1, imposing market clearing and solving for the general equilibrium results in the translog gravity equation:

$$\frac{x_{ij}/y_j}{n_i} = -\theta \ln(t_{ij}) + D_i + \theta \ln(T_j), \tag{1}$$

where x_{ij} is the bilateral trade flow between exporting country i and importing country j, y_j is the importer's income and n_i denotes the number of goods of country i (we ignore time indices for now). The dependent variable is the bilateral import share x_{ij}/y_j per good n_i of the exporting country. On the right-hand side, $\theta > 0$ is a translog preference parameter. Here D_i and T_j denote exporter and importer-specific terms given by:

$$D_i = \frac{y_i/y^W}{n_i} + \theta \sum_{s=1}^{S} \frac{y_s}{y^W} \ln\left(\frac{t_{is}}{T_s}\right)$$

[11] See Head and Mayer (2014) for an overview.

$$\ln(T_j) = \sum_{s=1}^{S} \frac{n_s}{N} \ln(t_{sj}),$$

where y^W denotes world income, S is the number of countries and N is the number of products in the world with $N \geq S$. Here T_j is akin to a multilateral resistance term since it represents a weighted average of bilateral trade costs.

The translog gravity equation (1) differs in two key respects from standard gravity equations as in Eaton and Kortum (2002) and Anderson and van Wincoop (2003). The dependent variable is the import share per good, which means that an empirical measure of n_i is required. In addition, the dependent variable is in levels. It is not the logarithmic bilateral trade flow. The gravity relationship is therefore not log-linear in trade costs. This implies a variable trade cost elasticity. This is the crucial feature we focus on in this paper.

More specifically, define the trade cost elasticity as $\eta \equiv \partial \ln(x_{ij})/\partial \ln(t_{ij})$. This is meant as the direct trade cost elasticity in the sense that indirect general equilibrium price effects are omitted here (we discuss those general equilibrium effects in detail in Online Appendix B.4). In standard gravity equations this elasticity would be constant.[12] In the translog gravity model, however, this elasticity is variable. It follows from (1) as:

$$\eta_{ij} = -\frac{\theta}{(x_{ij}/y_j)/n_i}. \tag{2}$$

That is, the trade cost elasticity is the preference parameter θ divided by the import share per good. Therefore, the larger a given import share, the smaller the trade cost elasticity in absolute magnitude. The ij subscript indicates that this elasticity varies by country pair.

In line with the literature, we assume that logarithmic trade costs $\ln(t_{ij})$ are a function of commonplace trade cost variables such as logarithmic bilateral distance, dummy variables for contiguity, common language as well as membership of trade agreements, currency unions and so on. As an example, let us consider a dummy variable for currency union membership CU_{ij} that takes on the value of one if countries i and j are both members, with coefficient κ. We expect κ to be negative since a currency union is generally thought to lower bilateral trade costs (our empirical results will confirm this). Based on the expression in (2), the effect of currency union membership on trade follows as:

$$\frac{\Delta \ln(x_{ij})}{\Delta CU_{ij}} \approx -\frac{\theta\kappa}{(x_{ij}/y_j)/n_i}, \tag{3}$$

where ΔCU_{ij} indicates entry into a currency union. Given that κ is likely negative, we expect a positive currency union effect on bilateral trade.

We would like to highlight a key aspect of our framework. As the denominators of (2) and (3) show, the heterogeneity of currency union effects is driven by variation across import shares. It would be conceivable that heterogeneity is instead driven by the currency union parameter in the trade cost function, perhaps because different currency unions have different trade cost effects.[13] However, as we show below in our empirical analysis, we find heterogeneous effects

[12] For instance, in Anderson and van Wincoop (2003) the elasticity would be equal to $1-\sigma$, where σ is the CES elasticity of substitution. In Eaton and Kortum (2002) it would be equal to the Fréchet shape parameter. In Chaney (2008) it would be equal to the Pareto shape parameter.

[13] Instead of the constant κ we would then have to adopt currency union-specific trade cost parameters. We allow for such an approach in our analysis of the euro in Subsection 2.1.4.

across different pairs *within* a given currency union, and even within a given bilateral country pair by direction of trade (due to bilaterally asymmetric import shares). This means that even if trade costs are bilaterally symmetric, trade cost effects can be bilaterally asymmetric in the translog gravity framework.[14]

In summary, the most important insight from this motivating framework is the variable trade cost effect in (2). Specifically, the trade cost effect should be larger in absolute magnitude for country pairs associated with smaller import shares. It also follows that a symmetric reduction in bilateral trade costs can lead to asymmetric increases in bilateral trade flows by direction of trade. These are testable predictions we now examine. While we also estimate the translog specification in (1) directly, we first turn towards an alternative approach.

2. Empirical Analysis

Our aim is to find out whether international trade data are characterised by variable trade cost elasticities. As a starting point, we first estimate gravity regressions with a standard constant trade cost elasticity. We then proceed by exploring variable trade cost elasticities. For that purpose, we adopt two approaches that are consistent with the theoretical framework in Section 1. The first approach is a modification of the standard gravity specification commonplace in the literature. Instead of estimating a single trade cost coefficient that is constant over the entire sample, we propose a flexible gravity framework that allows for heterogeneous trade cost effects across import shares. We explain this estimating strategy in more detail below (see Subsection 2.1). The second approach is to estimate the translog gravity equation (1) directly (see Subsection 2.2).

We introduce our approach with an application to the trade effect of currency unions. This is mainly for expositional purposes, and later in the paper we extend the analysis to other trade cost variables (see Section 3 in particular). We use a very large, comprehensive data set of aggregate annual bilateral trade flows that covers most of global trade in modern times. It consists of an unbalanced panel including 199 countries from 1949 to 2013. We provide details and descriptive statistics in Online Appendix A.

2.1. *Gravity with Heterogeneity*

This section describes our first approach. Focusing on currency unions, we start by estimating homogeneous trade cost effects with constant trade cost elasticities as in the standard gravity framework. We then modify the gravity specification and introduce a two-step procedure to allow for trade cost heterogeneity. We also explore the trade effect of the euro in more detail.

2.1.1. *Homogeneous trade cost effects in standard gravity*

As pointed out by Santos Silva and Tenreyro (2006), the validity of estimating a log-linear gravity model by OLS depends crucially on the assumption that the variance of the error term is independent from the regressors. Otherwise, the log transformation prevents the error term from having a zero conditional expectation, leading to inconsistent estimates of the true elasticities. PPML instead delivers consistent coefficient estimates, even in the presence of heteroskedasticity (Head and Mayer, 2014). Another advantage of PPML is that by expressing the dependent variable in levels, zero trade observations can be incorporated in the estimation (the log-linear gravity

[14] Trade flows can also be bilaterally asymmetric, but trade is balanced at the aggregate country level due to the general equilibrium nature of the model.

equation may suffer from selection bias as the zero values drop out of the regression). In our sample, around 35% of import shares are equal to zero (see Online Appendix A). In what follows we therefore estimate our regressions by PPML.[15]

We initially estimate homogeneous trade cost effects with a constant elasticity as in the standard gravity framework, based on PPML regressions. But we use the bilateral import share per good as the dependent variable to make sure our results are comparable to subsequent estimates that allow for heterogeneous trade cost effects. We thus estimate:

$$\left(\frac{x_{ij,t}/y_{j,t}}{n_{i,t}}\right) = \exp(\alpha_1 CU_{ij,t} + \alpha_2 Z_{ij,t} + D_{i,t} + D_{j,t} + D_{ij}) + \varrho_{ij,t}, \quad (4)$$

where we add time subscripts such that $x_{ij,t}$ is the bilateral Free on Board (FOB) export value from exporter i to importer j in year t, $y_{j,t}$ is country j's nominal GDP (both in US dollars) and $n_{i,t}$ denotes the number of goods in the exporting country that can be interpreted as an extensive margin measure. Trade costs depend on currency union membership $CU_{ij,t}$ that is a dummy variable taking on a value of one if countries i and j are both members in year t (and zero otherwise). Trade costs are also a function of other time-varying country pair variables $Z_{ij,t}$ that include dummy variables equal to one if both countries in the pair belong to an RTA or are members of the WTO, OECD and IMF in each year, and zero otherwise (Rose, 2005). We return to those other trade cost variables in Section 3.

To measure the exporting countries' extensive margin $n_{i,t}$, we collect each country's total exports by product category from United Nations Comtrade that are available from 1962 onwards. As the HS classification was only introduced in 1988, we rely on data at the four-digit HS level between 1988 and 2013, and at the four-digit SITC level from 1962 to 1987. We define the extensive margin as the number of different product categories exported by each country in each year relative to the total number of categories exported by all countries in the same year. Given that the Comtrade data are only available from 1962, have poor country coverage in some years and are reported according to two different classifications over time (i.e., SITC versus HS), we calculate the average extensive margin by exporter. This yields a time-invariant measure n_i, but we believe that this measure should provide us with useful information regarding the variation in the extensive margin across exporting countries.[16]

We later check the robustness of our findings by using an alternative proxy for the extensive margin (see Table C5 in Online Appendix C). We rely on the cross-country measure constructed by Hummels and Klenow (2005) using data on exports from 126 exporting to 59 importing countries in more than 5,000 six-digit HS-level product categories in 1995. We also assume that the extensive margin is unity for all exporters, in which case the dependent variable is simply the bilateral import share.

We control for an extensive set of fixed effects. We include time-varying exporter and importer fixed effects $D_{i,t}$ and $D_{j,t}$ to control for multilateral trade resistance and other exporter and importer-specific terms such as income. We also include country pair fixed effects D_{ij} to absorb

[15] We employ the `ppmlhdfe` Stata command written by Correia *et al.* (2020). It estimates a Poisson Pseudo Maximum Likelihood regression allowing for multiple levels of fixed effects. In the previous version of our paper (Chen and Novy, 2018) results were based on OLS estimation. Those results are generally the same qualitatively, although individual magnitudes may be different between PPML and OLS.

[16] For each year from 1962 to 1987, we calculate the number of four-digit SITC-level product categories exported by each country relative to the total number of four-digit SITC-level categories exported by all countries. For the years 1988 to 2013, we do the same at the four-digit HS level. For each country, the time-invariant n_i measure (that we use for our full sample between 1949 and 2013) is given by the average of the two series between 1962 and 2013.

all time-invariant bilateral trade frictions in each cross section. The country pair effects also help to control for the endogeneity of the currency union dummy if two countries deciding to join a currency union have traditionally traded a lot with each other (but the pair effects fail to do so if the two countries join following a surge in trade during the sample period; see Micco *et al.*, 2003 or Bun and Klaassen, 2007). Note that the pair effects are directional as non-directional pair effects would not capture the asymmetry in bilateral import shares within a pair. Identification is therefore achieved from the time series variation of each explanatory variable within a pair (e.g., from changes in bilateral currency union status over time).[17] To control for time-invariant idiosyncratic shocks correlated at the pair level (De Sousa, 2012), SEs are clustered at the non-directional country pair level. The coefficients to be estimated are denoted by the α. As sharing a common currency should be associated with more trade, we expect α_1 to be positive. The error term is $\varrho_{ij,t}$.

2.1.2. *Heterogeneous trade cost effects*

We then focus on trade cost heterogeneity. Our aim is to investigate whether the trade effect of currency unions, as captured by α_1 in specification (4), is heterogeneous across bilateral import shares per good, as predicted by the theoretical framework in Section 1. If we simply allowed α_1 to vary with import shares, we would have a simultaneity bias problem as the currency union effect would vary with the values taken by the dependent variable (Novy, 2013).

To address this issue, we modify the standard gravity specification by letting the currency union effect vary across *predicted* import shares. For that purpose, we propose a two-step methodology. In the first step we regress the import shares per good on geography-related variables (distance and contiguity) to generate the predicted shares. In the second step we investigate how the trade effect of currency unions varies across predicted shares. We now explain this approach in more detail.

In the first step we regress the import shares per good on time-invariant country pair controls and exporter-year and importer-year fixed effects:

$$\left(\frac{x_{ij,t}/y_{j,t}}{n_{i,t}}\right) = \exp(\delta K_{ij} + D_{i,t} + D_{j,t}) + \nu_{ij,t}. \tag{5}$$

Here K_{ij} includes geography-related variables, i.e., logarithmic bilateral distance and a contiguity dummy.[18] We do not include the time-varying pair variables for currency unions, RTAs, the WTO, OECD or IMF as they are not geography related and therefore more likely endogenous. We then generate the predicted shares that we denote by $((\widehat{x_{ij,t}/y_{j,t})}/n_{i,t})$.

In the second step we include an interaction term between the currency union dummy variable and the logarithmic predicted import shares, with ξ_2 as the key coefficient of interest.[19] We

[17] Baier *et al.* (2019) used directional pair effects to estimate the within-pair asymmetric effects of FTAs. The recent literature concludes that time-varying exporter and importer dummies and time-invariant country pair fixed effects should be included (De Nardis and Vicarelli, 2003; Baldwin, 2006; Baldwin and Taglioni, 2007; Baldwin *et al.*, 2008; Eicher and Henn, 2011; Mika and Zymek, 2018; Campbell and Chentsov, 2020). The earlier literature failed to do so (e.g., Rose, 2000).

[18] Our results remain robust if in the first-step regression (5) we allow for time-varying distance and contiguity coefficients, include further gravity variables or simply control for time-invariant (directional) country pair fixed effects (see Section 5). They also remain robust if we add the currency union, RTA, WTO, OECD and IMF variables.

[19] We interact with the logarithmic share since the dependent variable and the interaction term are then expressed in the same units, i.e., they are expressed as level shares if we exponentiate the logarithmic share on the right-hand side of (6).

estimate:

$$\left(\frac{x_{ij,t}/y_{j,t}}{n_{i,t}}\right) = \exp\left(\xi_1 CU_{ij,t} + \xi_2 CU_{ij,t} \times \ln\left(\widehat{\frac{x_{ij,t}/y_{j,t}}{n_{i,t}}}\right) + \xi_3 Z_{ij,t} + D_{i,t} + D_{j,t} + D_{ij}\right) + \varepsilon_{ij,t}. \quad (6)$$

Since this specification includes exporter-year, importer-year and country pair fixed effects, the main effect of the logarithmic predicted import share drops out of the regression. The trade effect of currency unions is given by $\xi_1 + \xi_2 \ln((x_{ij,t}\widehat{/y_{j,t})}/n_{i,t})$ and therefore depends on two components, i.e., the change in trade costs due to currency unions and the log predicted import share. If the trade effect of currency unions falls with bilateral import intensity as predicted by our theoretical framework, the interaction coefficient ξ_2 should be negative. As the log predicted import share is a generated regressor, we bootstrap SEs with a hundred replications.

The intuition of this two-step methodology is as follows. To avoid the simultaneity problem, the predicted import shares generated in the first step should not be correlated with the error term $\varepsilon_{ij,t}$ in the second step. The point of the first step is therefore to extract the exogenous component of import shares. We aim to achieve this by using geography-related regressors (distance and contiguity) as well as time-varying exporter and importer-specific fixed effects, and then predicting the import shares. In Online Appendix B we carry out Monte Carlo simulations to verify the validity of this two-step procedure. In Section 4 we explore the potential endogeneity of currency unions and reverse causality.

An alternative way of testing our prediction of heterogeneous currency union effects is to split the sample into intervals of predicted import shares per good ranked by value and to estimate:

$$\left(\frac{x_{ij,t}/y_{j,t}}{n_{i,t}}\right) = \exp(\beta_{1,h} CU_{ij,t} \times D_h + \beta_2 Z_{ij,t} + D_{i,t} + D_{j,t} + D_{ij} + D_h) + \epsilon_{ij,t}, \quad (7)$$

where D_h is a dummy variable for h equally sized intervals of predicted import shares per good. The currency union coefficient $\beta_{1,h}$ is estimated separately for each interval h. Consistent with our theoretical framework, we expect the currency union effect to be largest in the interval of lowest predicted shares, and to be weaker in intervals of higher shares.[20]

2.1.3. *Baseline results*

We start by discussing homogeneous currency union estimates. Before turning to the PPML estimation, in column (1) of Table 1 we first estimate the log-linear version of (4) by OLS, using the log bilateral import share per good as the dependent variable (the zero observations therefore drop out from the regression). The currency union coefficient is equal to 0.326, suggesting that a common currency is associated with an increase in bilateral trade of 39% on average ($\exp(0.326) - 1 = 0.385$). We note that this currency union estimate captures direct trade cost effects.[21] Joining an RTA, and becoming a member of the WTO, OECD and IMF are associated with an increase in bilateral trade (Rose, 2005).

[20] Quantile regressions could also be used to test our predictions. Various fixed effect estimators have recently been developed but little is known about their performance. Using the `qreg2` Stata command of Parente and Santos Silva (2016), we instead estimated pooled quantile regressions with clustered SEs. The currency union effect is overestimated due to the omission of the fixed effects, but we find that it falls with bilateral import shares.

[21] We discuss indirect general equilibrium effects in Online Appendix B.4. Those are second-order effects that are quantitatively small. The intuition is that currency unions are relatively rare at the bilateral level (see Online Appendix A), and they are only one out of several trade cost components. We also show that the general equilibrium effects are not systematically related to the heterogeneity of currency union effects.

Table 1. *Baseline Results.*

	(1)	(2)	(3)	(4)	(5)
CU	0.326***	0.202***	0.252***	0.964***	−0.382**
	(0.057)	(0.050)	(0.055)	(0.101)	(0.149)
CU × ln import share	–	–	–	0.283***	–
				(0.026)	
CU × ln predicted share	–	–	–	–	−0.197***
					(0.036)
RTA	0.415***	0.205***	0.127***	0.206***	0.123***
	(0.028)	(0.035)	(0.037)	(0.035)	(0.041)
WTO	0.146***	0.029	−0.004	0.031	−0.010
	(0.035)	(0.053)	(0.053)	(0.053)	(0.055)
OECD	0.366***	0.534***	0.590***	0.531***	0.582***
	(0.051)	(0.073)	(0.081)	(0.073)	(0.091)
IMF	0.165**	0.321***	0.203*	0.334***	0.203*
	(0.065)	(0.106)	(0.105)	(0.106)	(0.105)
CU estimates					
Mean	–	–	–	−1.356***	0.980***
				(0.138)	(0.129)
10th percentile	–	–	–	−2.573***	1.505***
				(0.244)	(0.218)
90th percentile	–	–	–	−0.299***	0.478***
				(0.062)	(0.065)
Observations	780,818	780,818	1,131,641	780,818	1,131,641
Zeros included	No	No	Yes	No	Yes
Estimator	OLS	PPML	PPML	PPML	PPML

Notes: Exporter-year, importer-year and (directional) country pair fixed effects are included. Robust SEs adjusted for clustering at the (non-directional) country pair level are reported in parentheses in (1) to (4). SEs are bootstrapped in (5). ***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively. The dependent variable is the import share per good but the log import share per good in (1). 'predicted share' is the predicted import share per good.

Next, we regress (4) by PPML but we omit the zero observations from the sample. This allows us to assess how PPML affects coefficient estimates. As explained by Mayer *et al.* (2019), OLS and PPML estimates may differ because the two estimators have different first-order conditions (Eaton *et al.*, 2013; Head and Mayer, 2014). While the OLS conditions involve logarithmic deviations of trade from its expected value, the PPML conditions involve level deviations and therefore tend to give more weight to country pairs with high levels of trade.[22] If country pairs with higher import intensity have smaller trade cost elasticities, as we argue, then the PPML currency union coefficient should be smaller than its OLS counterpart.

In column (2) the currency union estimate decreases to 0.202 such that sharing a common currency is associated with 22% more trade on average. This result is therefore consistent with our prediction of country pairs with higher import shares having a smaller trade cost elasticity.[23] It is also consistent with papers finding that PPML reduces the size and significance of currency union estimates (Santos Silva and Tenreyro, 2010; Mika and Zymek, 2018; Larch *et al.*, 2019; Mayer *et al.*, 2019). The WTO estimate becomes insignificant, while the RTA, OECD and IMF

[22] To illustrate that PPML gives more weight to the country pairs with high levels of trade, Mayer *et al.* (2019) followed the approach of Eaton *et al.* (2013), who estimated a multinomial gravity model for aggregate bilateral trade shares. They regressed the trade share (bilateral trade divided by total trade) by PPML, and as trade shares give less weight to the large trade flows in levels, they found that the coefficient estimates lie between the OLS and PPML estimates of regressing bilateral trade flows.

[23] By finding that the PPML estimates for RTAs and sharing the euro are smaller than their OLS counterparts, Mayer *et al.* (2019) concluded that country pairs with high volumes of trade have a smaller trade cost elasticity.

coefficients remain positive (the RTA coefficient is smaller, while the OECD and IMF coefficients are larger than their OLS counterparts).[24]

In column (3) we include the zero observations in the sample. Compared to column (2), the currency union coefficient increases but only slightly to 0.252, suggesting that a common currency is associated with 29% more trade on average. The coefficients on the other time-varying pair controls do not change much either. Consistent with the literature, we therefore confirm that including the zero observations in the sample does not substantially affect coefficient estimates (see, for instance, Mika and Zymek, 2018 or Mayer *et al.*, 2019).

Our next task is to demonstrate whether these results mask heterogeneity in the trade effect of currency unions across country pairs. Purely as an illustration, in column (4) we interact the currency union dummy with *actual* logarithmic import shares per good. The coefficient on the interaction term is strongly positive. This is driven by the fact that, for all currency union pairs, the interaction term contains the same values as the dependent variable. To be clear, this specification is misguided as it suffers from simultaneity bias. We do not recommend it, and we only include it here for comparative purposes.

To address the simultaneity bias, we proceed with our two-step methodology. We estimate the first-step regression (5). Import intensity is stronger between less distant and contiguous countries (the estimated coefficients are significant at the 1% level). We then generate the log predicted import shares per good and estimate the second-step regression (6). Column (5) shows that the coefficient on the interaction term is negative and significant. The impact of currency unions is thus heterogeneous as it *falls* with predicted import shares.[25] It is clear that the two-step approach counteracts the simultaneity bias in column (4) as the sign of the interaction coefficient flips (also see Online Appendix B).[26]

In the lower part of Table 1 we report the implied currency union estimates at the mean and different percentiles of the log predicted import share distribution as well as the corresponding SEs. When we interact the currency union dummy with actual log import shares in column (4), due to the simultaneity bias those estimates erroneously suggest that currency unions are associated with smaller import shares.[27] But once we interact with the log predicted import shares in column (5), the magnitude of the currency union estimate is positive and, most importantly, it goes down when we move from the 10th to the 90th percentile (i.e., from small to large shares). Specifically, while the currency union estimate is 0.980 at the mean value of log predicted shares, it is 1.505 for a country pair at the 10th percentile and 0.478 at the 90th percentile.[28] In other words, at the 10th percentile currency unions are associated with 350% higher import shares ($\exp(1.505) - 1 = 3.504$), whereas at the 90th percentile the corresponding effect is only 61%.

To shed further light on the heterogeneity, we choose a few examples of country pairs from across the world with either small or large import shares in the final year of our sample. Based on the estimates of column (5) in Table 1, we report the associated currency union estimates in

[24] See Section 3 for a discussion of the magnitude of these coefficients.

[25] Irarrazabal *et al.* (2015) introduced additive trade costs. Under a broad range of demand systems, additive trade costs work to reduce the elasticity in magnitude. That is, ceteris paribus bilateral pairs with larger additive costs and thus a smaller trade share are associated with weaker (not stronger) elasticities.

[26] If we omit the zero import shares from the sample and estimate the log-linear versions of (5) and (6) by OLS, the trade effect of currency unions also falls with predicted import shares. See the previous version of our paper (Chen and Novy, 2018).

[27] See Online Appendix B.2 where we provide Monte Carlo simulation results on this point.

[28] The elasticities at the mean, the 10th and 90th percentiles are calculated for non-zero import shares. The currency union estimate at the mean of log predicted shares underestimates the elasticity at the 10th percentile by 53% (1.505/0.980), and overestimates the elasticity at the 90th percentile by 51% (0.478/0.980).

Table 2. *Examples of Pair-Specific Currency Union Effects.*

Large effects			Small effects		
Exporter	Importer	CU estimates	Exporter	Importer	CU estimates
Niger	Equatorial Guinea	1.360*** (0.192)	India	Bhutan	0.080 (0.080)
Central African Republic	Mali	1.321*** (0.186)	Senegal	Guinea-Bissau	0.001 (0.090)
Liberia	The Bahamas	1.142*** (0.155)	Spain	Portugal	−0.012 (0.092)
Chad	Côte d'Ivoire	0.878*** (0.113)	Germany	Netherlands	−0.109 (0.106)

Notes: The currency union estimates for each country pair are calculated based on the estimates of column (5) in Table 1. They are evaluated at log predicted import shares for the year 2013. Bootstrapped SEs adjusted for clustering at the (non-directional) country pair level are reported in parentheses. *** indicates significance at the 1% level.

Table 3. *Examples of Pair-Specific Bilateral Asymmetries.*

				Actual data CU = 1		Counterfactual CU = 0	
Exporter	Importer	Import share	CU estimates	Bilateral trade	Bilateral balance	Bilateral trade	Bilateral balance
Austria	Germany	1.58%	0.481*** (0.066)	\$57.1bn	−\$25.8bn	\$21.8bn	−\$58.4bn
Germany	Austria	20.07%	0.032 (0.086)	\$82.9bn	\$25.8bn	\$80.2bn	\$58.4bn
Belgium/Lux.	France	2.96%	0.303*** (0.063)	\$79.6bn	\$30.4bn	\$51.4bn	\$2.4bn
France	Belgium/Lux.	8.79%	0.005 (0.090)	\$49.3bn	−\$30.4bn	\$49.0bn	−\$2.4bn
Togo	Côte d'Ivoire	0.05%	0.497*** (0.067)	\$15.2m	−\$97.9m	\$5.4m	−\$83.0m
Côte d'Ivoire	Togo	2.60%	0.197*** (0.069)	\$113.1m	\$97.9m	\$88.5m	\$83.0m
Spain	Netherlands	1.30%	0.460*** (0.064)	\$10.2bn	−\$9.4bn	\$4.2bn	−\$3.8bn
Netherlands	Spain	1.45%	0.465*** (0.065)	\$19.6bn	\$9.4bn	\$8.0bn	\$3.8bn

Notes: The currency union estimates for each country pair are calculated based on the estimates of column (5) in Table 1. They are evaluated at log predicted import shares for the year 2013. Bootstrapped SEs adjusted for clustering at the (non-directional) country pair level are reported in parentheses. *** indicates significance at the 1% level. The bilateral trade data for CU = 1 in 2013 are calculated by applying the growth rates of exports reported by the International Monetary Fund's Direction of Trade Statistics to the bilateral exports data provided by Head *et al.* (2010). See Online Appendix A for details.

Table 2 (evaluated at log predicted import shares for the year 2013). Currency union effects are large for country pairs with small import shares. Naturally, these include thin trading relationships such as Equatorial Guinea importing from Niger (290%), Mali from the Central African Republic (275%), the Bahamas from Liberia (213%) and Côte d'Ivoire from Chad (141%). Conversely, some country pairs with large import shares do not tend to be associated with increased trade shares through currency unions, the effect being insignificant for Bhutan importing from India, Guinea-Bissau from Senegal, Portugal from Spain and the Netherlands from Germany.

We also find that the trade effect of currency unions is heterogeneous *within country pairs* and therefore asymmetric by direction of trade. In Table 3 we provide examples of country pairs with bilateral asymmetries in currency union effects (evaluated at log predicted import shares for the year 2013). For instance, the effect is large (62%) when Germany imports from Austria (a low share), but small (3%) and insignificant when Austria imports from Germany (a high share). The effect is also relatively large when France imports from Belgium-Luxembourg and when Côte d'Ivoire imports from Togo (low shares) but insignificant or small in the other direction (high shares). By contrast, as Spain and the Netherlands have similar bilateral import shares, using the same currency is associated with a similar effect in either direction.

Let us consider the example of Germany and Austria in more detail. The import share of Germany from Austria is low at 1.6% for the year 2013. But in the opposite direction the import share is large at 20.1%.[29] The corresponding trade flow values are \$57.1bn and \$82.9bn (i.e., Austria has a bilateral trade deficit with Germany). As a counterfactual exercise, suppose that these two countries were no longer in a currency union, i.e., their bilateral currency union dummy would switch from 1 to 0, and bilateral trade costs would go up. According to (3) and the estimates in Table 3, the import share from Austria to Germany would decrease by 62% ($\exp(0.481) - 1 = 0.618$) all else being equal. In the data this would imply a reduction of trade from Austria to Germany by about \$35.3bn (to \$21.8bn). The import share in the other direction would only decrease by 3%, corresponding to a reduction of trade from Germany to Austria by roughly \$2.7bn (to \$80.2bn). Thus, the bilateral trade deficit would widen (from \$25.8bn to \$58.4bn). Vice versa, a reduction in bilateral trade costs would shrink the bilateral trade deficit in this particular case. We note that the translog gravity equation (1) is consistent with bilateral imbalances even if bilateral trade costs are symmetric as in our example, although in the theoretical model aggregate trade remains balanced through general equilibrium adjustments.[30] Table 3 also reports the results of the corresponding counterfactual exercise for the other country pairs in the table.

Given our reliance on the two-step procedure outlined above, we check in detail whether our methodology is valid and does not lead to spurious heterogeneity. In Online Appendix B we carry out Monte Carlo simulations to verify the validity of our procedure. When we assume that the data generating process is driven by variable trade cost elasticities as in our theoretical framework in Section 1, our regressions based on the two-step procedure indeed produce heterogeneous effects that match those implied by the model underlying the data generating process, both qualitatively and quantitatively.[31] Conversely, if standard gravity were the data generating process, we demonstrate that our two-step procedure would not spuriously produce heterogeneous trade cost effects.

We proceed by regressing (7). In Table 4 we report currency union effects estimated separately by intervals of log predicted import shares per good. Based on the median of log predicted shares, column (1) of Table 4 splits the data into two intervals where the first interval includes the lower shares. As expected, the currency union coefficient is larger (equal to 0.736) for the lower shares and smaller (equal to 0.238) for the larger shares. Column (2) splits the sample into three equally sized intervals of log predicted import shares per good. The magnitude of the currency union coefficient declines from the first to the last interval. In column (3) we split the data into three intervals but in such a way that each includes the same number of observations for which the

[29] While the currency union estimates in Table 3 are based on import shares per good, for simplicity, we frame our example in terms of import shares. In our data the bilateral import shares for Germany and Austria (1.6% and 20.1%) are roughly the same as our measure of bilateral import shares per good (1.9% and 20.5%).

[30] Apart from the direct trade cost effect mentioned in the text, indirect price index effects would also be in operation (see Online Appendix B.4).

[31] Also see Figure 1 in Section 4.

Table 4. *Heterogeneous Currency Union Effects: Intervals.*

	(1)	(2)	(3)
CU (first interval)	0.736***	0.802***	0.894***
	(0.173)	(0.214)	(0.187)
CU (second interval)	0.238***	0.668***	0.721***
	(0.054)	(0.154)	(0.121)
CU (third interval)	–	0.224***	0.205***
		(0.054)	(0.054)
RTA	0.126***	0.125***	0.128***
	(0.037)	(0.037)	(0.037)
WTO	−0.004	−0.004	−0.007
	(0.053)	(0.053)	(0.053)
OECD	0.588***	0.585***	0.586***
	(0.081)	(0.081)	(0.081)
IMF	0.205*	0.204*	0.203*
	(0.105)	(0.105)	(0.105)
Intervals split by	# obs.	# obs.	# obs. CU = 1
Observations	1,131,641	1,131,641	1,131,641

Notes: PPML estimation. Exporter-year, importer-year, (directional) country pair and interval fixed effects are included. Robust SEs adjusted for clustering at the (non-directional) country pair level are reported in parentheses. *** and * indicate significance at the 1% and 10% levels, respectively. The dependent variable is the import share per good.

currency union dummy is equal to one. As before, the magnitude of the currency union estimate falls with predicted shares (in all columns, we can reject at the 1% level that the coefficients are equal across intervals).[32]

2.1.4. *The euro*

Given the prominence of the European single currency, we investigate the trade effect of the euro in more detail. In column (1) of Table 5 we first estimate specification (4) but the currency union dummy is split between euro and non-euro currencies. Sharing a common currency is associated with 18% more trade for the euro ($\exp(0.163) - 1 = 0.177$), and 36% more trade for non-euro currencies ($\exp(0.309) - 1 = 0.362$). In column (2) we interact the currency union indicators with log predicted import shares, and we observe heterogeneity in the trade effects of both euro and non-euro currency unions.

As argued by previous authors, one issue with the regressions in columns (1) and (2) is that they fail to control for the effect of European integration more broadly. As a result, the trade impact of the euro is likely to be overestimated because it confounds the effect of European integration with the effect of the single currency (see Baldwin, 2006 for a discussion). To address this issue, in column (3) we further include a time trend for EU countries (both in and out of the euro) to control for the ongoing European integration process (Micco *et al.*, 2003; Baldwin, 2006; Bun and Klaassen, 2007; Baldwin *et al.*, 2008; Berger and Nitsch, 2008; Mika and Zymek, 2018; Campbell and Chentsov, 2020). The positive coefficient on the trend indicates that, on average, EU countries trade more intensively with each other over time.[33] Still, the inclusion of the trend

[32] In column (2) the number of observations for which the currency union dummy is equal to one is 3,404 in the first interval, 6,227 in the second and 8,225 in the third. In column (3) the number of observations for which the currency union dummy is equal to one is 5,952 in each of the three intervals.

[33] We include a trend for EU countries as EU integration has affected all EU countries whether or not they have adopted the euro (Baldwin, 2006). The trend controls for EU policies such as the Single Market, treaties on EU integration, the Exchange Rate Mechanism, etc. It is included for twenty-seven EU countries (as Belgium and Luxembourg are merged together) and for the EU overseas territories. The trend varies across country pairs as it only starts in the year once the two countries in a pair are both members of the EU. Our results remain similar if we do not include a trend for the overseas

Table 5. *The Euro.*

	(1)	(2)		(3)	
CU non-euro	0.309***	−0.465**		−0.396**	
	(0.080)	(0.207)		(0.201)	
CU non-euro × ln predicted share	–	−0.234***		−0.217***	
		(0.048)		(0.046)	
Euro	0.163**	−0.043		−0.564***	
	(0.067)	(0.167)		(0.151)	
Euro × ln predicted share	–	−0.068*		−0.119***	
		(0.040)		(0.036)	
RTA	0.129***	0.126***		0.112***	
	(0.037)	(0.041)		(0.041)	
WTO	−0.008	−0.013		0.009	
	(0.053)	(0.055)		(0.055)	
OECD	0.591***	0.581***		0.539***	
	(0.081)	(0.091)		(0.088)	
IMF	0.200*	0.202*		0.223**	
	(0.105)	(0.105)		(0.105)	
Trend EU countries	–	–		0.026***	
				(0.003)	
CU estimates		Non-euro	Euro	Non-euro	Euro
Mean	–	1.151***	0.425***	1.108***	0.259**
		(0.172)	(0.147)	(0.169)	(0.130)
10th percentile	–	1.774***	0.605**	1.687***	0.576***
		(0.286)	(0.246)	(0.281)	(0.218)
90th percentile	–	0.555***	0.253***	0.553***	−0.045
		(0.096)	(0.078)	(0.096)	(0.069)
Observations	1,131,641	1,131,641		1,131,641	

Notes: PPML estimation. Exporter-year, importer-year and (directional) country pair fixed effects are included. Robust SEs adjusted for clustering at the (non-directional) country pair level are reported in parentheses in (1). SEs are bootstrapped in (2) and (3). ***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively. The dependent variable is the import share per good. 'predicted share' is the predicted import share per good.

does not affect our main insight as both euro and non-euro currency unions are associated with heterogeneous trade effects.

As shown in the lower part of Table 5, at the mean, the 10th and 90th percentiles of log predicted shares the euro estimates are smaller in magnitude once we include the trend (in column (3)).[34] They are generally weaker than the estimates for non-euro currency unions. As the average import share per good in our sample is significantly larger for euro member pairs compared to non-euro currency union pairs (they are equal to 2% and 1.4%, respectively), the finding that the euro trade effect is weaker on average is thus consistent with the theoretical framework in Section 1. It is also consistent with evidence showing that PPML reduces the size and significance of the euro trade effect (Santos Silva and Tenreyro, 2010; Mika and Zymek, 2018; Larch *et al.*, 2019; Mayer *et al.*, 2019).

In column (3) the euro estimate is equal to 0.576 for a country pair at the 10th percentile of log predicted shares. Examples of euro country pairs with small import shares associated with large trade effects are Ireland importing from Cyprus (70%), Finland from Malta (50%) and Finland from Greece (23%). In contrast, euro country pairs with large import shares that are

territories. They also remain similar if we interact the trend with country pair dummy variables, but in that case we were unable to bootstrap the SEs as the resampled samples encountered issues in fitting our model and the replications failed to converge.

[34] Bun and Klaassen (2007), Berger and Nitsch (2008) and Mika and Zymek (2018) also found that the inclusion of a time trend reduces the magnitude of the euro trade effect.

Table 6. *Translog Estimation.*

	(1)	(2)
CU	0.006***	0.003***
	(0.001)	(0.001)
RTA	0.005***	0.004***
	(0.000)	(0.000)
WTO	−0.001	0.000
	(0.000)	(0.000)
OECD	0.003***	0.002***
	(0.001)	(0.001)
IMF	0.001**	0.000
	(0.001)	(0.000)
CU estimates		
Mean	0.912***	0.470***
	(0.235)	(0.137)
10th percentile	1,514.591***	780.223***
	(391.203)	(227.256)
90th percentile	0.484***	0.250***
	(0.125)	(0.073)
Observations	780,818	1,203,322
Zeros included	No	Yes
R^2	0.644	0.588

Notes: OLS estimation. Exporter-year, importer-year and (directional) country pair fixed effects are included. Robust SEs adjusted for clustering at the (non-directional) country pair level are reported in parentheses. *** and ** indicate significance at the 1% and 5% levels, respectively. The dependent variable is the import share per good.

not associated with increased trade shares include Cyprus importing from Greece and Germany importing from Italy (the effects are insignificant). We also find evidence of heterogeneity by direction of trade. For instance, the trade effect of sharing the euro is large when Austria imports from Malta (low predicted shares). But it is insignificant when Malta imports from Austria (high predicted shares).

Although our primary objective is not to determine whether the effect of the euro is stronger or weaker on average compared to other currency unions, our results suggest that its effect on trade is more modest compared to other common currencies. Yet our main interest is in the heterogeneity of trade cost effects. Consistent with the predictions of our model we find that the euro effect is heterogeneous across and within country pairs.[35]

2.2. *Translog Approach*

We now report the results of implementing our second approach where we estimate the translog gravity equation (1) directly using OLS estimation. We can then compute the pair-specific currency union effects with the help of (3). We report two sets of results, i.e., excluding and including the zero import share observations in the sample. We note that in contrast to the two-step regressions reported earlier, the translog approach does not require us to predict the bilateral import shares per good in a first step.

Column (1) of Table 6 reports the results excluding the zero observations from the sample and the currency union coefficient is equal to 0.006. As shown in the lower part of the table,

[35] In Table C2 of Online Appendix C we show that our main results are robust to controlling for wars and political conflicts that have been argued to drive the trade effect of currency unions (Campbell, 2013; Campbell and Chentsov, 2020). As the EU has essentially experienced a period of uninterrupted peace since the end of World War II, our results for the euro provide further evidence that our findings are not driven by geopolitical events.

this corresponds to an estimate of 0.912 at the mean value of import shares. As predicted by the translog framework, the effect is heterogeneous across country pairs, and the currency union estimate decreases from the 10th to the 90th percentile of import shares per good. However, we note that the currency union estimate at the 10th percentile is extremely large compared to previous tables. The reason is that translog imposes a hyperbolic functional form for the calculation of trade cost elasticities. This can be seen in (3) in that the estimated coefficient, $-\theta\kappa$, is divided by import shares. Since import shares at low percentiles are very close to zero (see the descriptive statistics in Online Appendix A), the implied elasticities mechanically become very large. We therefore treat the currency union estimates at low percentiles with a particular degree of caution. Besides, all other regressors are significant with the expected signs, with the exception of the WTO dummy variable.

In column (2) when we include the zero observations in the sample, the heterogeneity across percentiles of import shares continues to hold. That is, sharing a common currency is associated with more bilateral trade, and this effect is stronger for country pairs with smaller import shares. The magnitude of the currency union estimate at the mean value of (non-zero) import shares is smaller at 0.470. This magnitude corresponds to the mean estimate of 0.980 in column (5) of Table 1. The mean effect is thus larger in Table 1 but due to the specific translog functional form the heterogeneity in Table 6 is more pronounced.

3. Other Trade Cost Variables

In the previous section we focused on currency union effects but this was mainly for exposition. Our approach is equally applicable to other trade cost components, and we discuss them now in turn. We discuss trade cost components represented by dummy variables (such as membership of trade agreements) that already appeared in earlier regression tables, and we also discuss continuous trade cost variables such as bilateral distance and tariffs.

We present the regression results in Table 7. In column (1) we estimate (6) where we interact the currency union dummy with log predicted import shares, but in the same way we now also interact the other time-varying trade cost components represented by dummy variables. The coefficients on the RTA and WTO interaction terms are negative. The trade effects of RTAs and the WTO are thus heterogeneous and smaller for country pairs with larger import shares. Specifically, the coefficient on the RTA dummy variable is equal to 0.489 at the mean value of log predicted shares, 0.773 for a country pair at the 10th percentile and 0.218 at the 90th percentile. Joining an RTA is thus associated with 117% more bilateral trade ($\exp(0.773) - 1 = 1.166$) at the 10th percentile, and 24% more trade only at the 90th percentile. The coefficient on the WTO dummy variable is equal to 0.137 at the mean value of log predicted shares, 0.295 at the 10th percentile and -0.014 at the 90th percentile (the latter is insignificant). Becoming a member of the WTO is thus associated with 34% more bilateral trade at the 10th percentile, and with no change in trade at the 90th percentile. These findings are broadly consistent with the literature showing that the trade effects of trade agreements and WTO membership are heterogeneous.[36]

The coefficients on the IMF and OECD interaction terms are positive, however. The trade effects of joining these organisations therefore increase with bilateral import shares. This finding is consistent with columns (1) and (2) of Table 1 that show that the PPML coefficients for the

[36] Key references that stress heterogeneity include Glick (2017) and Baier *et al.* (2019) for RTAs, and Subramanian and Wei (2007) and Felbermayr *et al.* (2020) for WTO membership.

Table 7. *Heterogeneous Trade Cost Effects.*

	(1)	(2)	(3)	(4)	(5)	(6)
CU	−0.367**	−0.831***	−0.614***	−0.061	−0.434***	−0.750***
	(0.158)	(0.232)	(0.218)	(0.042)	(0.122)	(0.229)
CU × ln predicted share	−0.190***	−0.354***	−0.298***	–	−0.131***	−0.209***
	(0.037)	(0.049)	(0.046)		(0.031)	(0.050)
RTA	−0.246**	0.195	0.504***	−0.058	−0.545***	0.421**
	(0.113)	(0.158)	(0.174)	(0.041)	(0.122)	(0.213)
RTA × ln predicted share	−0.106***	−0.042	0.042		−0.168***	0.029
	(0.025)	(0.032)	(0.036)		(0.028)	(0.043)
WTO	−0.273***	−0.107	−0.081	0.172***	0.107	−0.204
	(0.096)	(0.131)	(0.132)	(0.039)	(0.140)	(0.214)
WTO × ln predicted share	−0.059***	−0.039*	−0.036	–	−0.044	−0.141***
	(0.017)	(0.022)	(0.023)		(0.029)	(0.038)
OECD	1.047***	−1.127***	−1.039***	0.138**	0.386	−0.225
	(0.202)	(0.211)	(0.216)	(0.062)	(0.255)	(0.203)
OECD × ln predicted share	0.111***	−0.144***	−0.125***	–	0.018	−0.021
	(0.034)	(0.042)	(0.043)		(0.046)	(0.039)
IMF	0.537***	0.550***	0.611***	0.324***	0.465	0.409
	(0.151)	(0.188)	(0.176)	(0.115)	(0.524)	(0.551)
IMF × ln predicted share	0.077***	0.076***	0.073***	–	−0.041	0.007
	(0.023)	(0.023)	(0.023)		(0.069)	(0.055)
ln Distance	–	−0.990***	−0.381	–	–	−0.595**
		(0.036)	(0.236)			(0.294)
ln Distance × ln predicted share	–	–	0.048***	–	–	0.039*
			(0.017)			(0.021)
Contiguity	–	0.505***	0.017	–	–	−0.001
		(0.071)	(0.172)			(0.202)
Contiguity × ln predicted share	–	–	−0.094*	–	–	−0.169***
			(0.049)			(0.050)
Shared language	–	0.513***	0.347**	–	–	0.263
		(0.056)	(0.144)			(0.161)
Shared language × ln predicted share	–	–	−0.042	–	–	−0.039
			(0.028)			(0.031)
ln (1+tariff)	–	–	–	−0.380**	–	–
				(0.176)		
ln (1+tariff) × ln predicted share	–	–	–	–	0.190**	0.471***
					(0.086)	(0.161)
ln (exporter GDP × importer GDP)	–	–	–	0.266***	–	–
				(0.031)		
Observations	1,131,641	1,131,641	1,131,641	356,491	368,733	388,798
Exporter-year fixed effects	Yes	Yes	Yes	No	Yes	Yes
Importer-year fixed effects	Yes	Yes	Yes	No	Yes	Yes
Pair fixed effects (directional)	Yes	No	No	Yes	Yes	No
Exporter fixed effects	No	No	No	Yes	No	No
Importer fixed effects	No	No	No	Yes	No	No
Year fixed effects	No	No	No	Yes	No	No

Notes: PPML estimation. Bootstrapped SEs adjusted for clustering at the (non-directional) country pair level are reported in parentheses. ***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively. The dependent variable is the import share per good. 'predicted share' is the predicted import share per good. The weighted mean applied tariff rate (in percentage terms) only varies by importer and year.

IMF and OECD dummy variables are larger than their OLS counterparts. As PPML gives more weight to country pairs with high levels of trade, the larger coefficients indicate that the effects of IMF and OECD membership are stronger for country pairs that trade intensively.[37] But we

[37] By contrast, columns (1) and (2) of Table 1 show that the PPML estimates for currency unions, RTAs and the WTO are smaller than their OLS counterparts. This is consistent with column (1) of Table 7 that shows that these variables are associated with smaller increases in bilateral trade for country pairs with large import shares.

believe that these findings should be interpreted with caution because it is not clear to what extent the purpose of the two organisations is focused on the reduction of trade costs. As explained by Rose (2005), although the IMF and the OECD are interested in trade creation, they also have various objectives other than trade promotion. This contrasts with the WTO that is primarily concerned with trade liberalisation and arguably can more easily be viewed as having a trade cost reducing effect.

In column (2) of Table 7 we add bilateral distance and dummy variables for sharing a common border and a common language and therefore drop the country pair fixed effects.[38] In column (3) we further interact these three controls with log predicted import shares. The trade effects of all three variables are strongly heterogeneous across bilateral import shares (Novy, 2013). At the mean value of log predicted shares, distance reduces trade with a coefficient of −0.711, while sharing a common border and a common language promotes trade with coefficients of 0.668 and 0.639. These effects are larger in magnitude at the 10th percentile of log predicted shares (the coefficients on distance, contiguity and common language are −0.838, 0.919 and 0.752). They are smaller at the 90th percentile (the coefficients are −0.589, 0.428 and 0.532). All the while, the results for currency unions remain. But the omission of country pair fixed effects in columns (2) and (3) turns the RTA interaction term insignificant. This suggests that the trade effects of RTAs are heterogeneous when countries join an RTA but not necessarily in the cross section.

Next, we investigate whether tariffs have a heterogeneous effect on import shares. From the World Bank's World Development Indicators we extract the weighted mean effectively applied tariff rate (in percentage terms) imposed by each importing country on all products from all trading partners.[39] As the data are only available from 1988, our sample size is significantly reduced. Also, as the weighted mean tariff is specific to each importing country and is not defined on a bilateral basis, it simply captures each country's overall degree of protectionism.

To get a sense of the homogeneous effect of tariffs on bilateral import shares, in column (4) of Table 7 we regress the import shares per good on dummy variables for currency unions, RTA, WTO, OECD and IMF membership, the logarithm of one plus the tariff rate, the logarithmic product of exporter and importer GDP, as well as year fixed effects and time-invariant exporter, importer and country pair fixed effects. On average, the currency union and RTA effects are insignificant, while WTO, OECD and IMF membership are associated with more trade. As expected, tariffs are associated with reduced bilateral import shares. But we stress that this specification does not include time-varying exporter and importer fixed effects and should be interpreted with caution.

In column (5), we estimate the same specification as in column (1) but include an interaction term between tariffs and log predicted import shares. As we now include time-varying importer fixed effects, the main effect of tariffs drops out but the coefficient on the tariff interaction term is positive and significant. The (negative) effect of tariffs is therefore heterogeneous and smaller in magnitude for the country pairs with larger import shares. The effects of currency unions and RTAs also fall with bilateral import shares. The WTO, IMF and OECD variables and their interaction terms with log predicted shares are insignificant.

[38] As there are no country pair fixed effects, the sample size is slightly larger as fewer singletons perfectly predicted by the fixed effects are dropped.

[39] The effectively applied bilateral tariffs are preferential rates if applicable, and Most Favoured Nation ones otherwise. For each importer, they are averaged across bilateral partners using product import shares as weights. In results available upon request we show that our results remain similar if we instead use the weighted Most Favoured Nation tariff rates of each country on all imports from the rest of the world.

Table 8. *Heterogeneous Trade Cost Effects (One by One).*

	(1)	(2)	(3)	(4)	(5)
CU	-0.387^{**}	–	–	–	–
	(0.151)				
CU × ln predicted share	-0.202^{***}	–	–	–	–
	(0.036)				
RTA	–	-0.191^{*}	–	–	–
		(0.112)			
RTA × ln predicted share	–	-0.099^{***}	–	–	–
		(0.026)			
WTO	–	–	-0.152	–	–
			(0.093)		
WTO × ln predicted share	–	–	-0.033^{**}	–	–
			(0.016)		
OECD	–	–	–	0.945^{***}	–
				(0.181)	
OECD × ln predicted share	–	–	–	0.074^{**}	–
				(0.032)	
IMF	–	–	–	–	0.511^{***}
					(0.148)
IMF × ln predicted share	–	–	–	–	0.070^{***}
					(0.021)
Trade cost estimates					
Mean	1.008^{***}	0.497^{***}	0.074	0.435^{***}	0.027
	(0.129)	(0.081)	(0.065)	(0.105)	(0.115)
10th percentile	1.545^{***}	0.761^{***}	0.162^{*}	0.239	-0.160
	(0.218)	(0.145)	(0.095)	(0.170)	(0.145)
90th percentile	0.493^{***}	0.243^{***}	-0.009	0.623^{***}	0.205^{*}
	(0.067)	(0.038)	(0.055)	(0.088)	(0.108)
Observations	1,131,641	1,131,641	1,131,641	1,131,641	1,131,641

Notes: PPML estimation. Exporter-year, importer-year and (directional) country pair fixed effects are included. Bootstrapped SEs adjusted for clustering at the (non-directional) country pair level are reported in parentheses. ***, ** and * indicate significance at the 1%, 5% and 10% levels, respectively. The dependent variable is the import share per good. 'predicted share' is the predicted import share per good.

In column (6) we omit country pair fixed effects and add distance, dummy variables for sharing a common border and a common language, and their interactions with log predicted import shares. Our results of heterogeneous trade cost effects continue to hold (although the interaction terms for RTAs and a common language become insignificant in this particular specification).

Finally, in Table 8 we run simple specifications where we only include the individual time-varying trade cost dummy variables and their interaction terms, one at a time. As before, we report the implied trade cost estimates at the bottom of the table. The coefficient patterns as well as the trade cost estimates are similar to the previous table.

Overall, we conclude that the predictions of our model are not limited to the effects of currency unions. They apply more generally to a large set of popular trade cost variables including RTAs, WTO membership, bilateral distance, sharing a common border, a common language and tariffs.

4. Simulation and Endogeneity

Many popular trade cost variables in the literature are potentially endogenous. For example, it would be implausible to assume that the formation of trade agreements or currency unions is exogenous to countries' trading patterns. Our aim is to explore the effect of trade cost endogeneity on our preferred two-step estimation procedure. For that purpose, we run simulations and demonstrate the validity of the procedure. We again lean on currency unions as the main illustration.

Currency unions are not randomly assigned. Santos Silva and Tenreyro (2010) argued that joining a currency union becomes more likely when countries are geographically close, speak the same language and have a former colonial link. Persson (2001) addressed selection on observables. By applying a propensity-score matching estimator, he accounted for the fact that characteristics such as distance and trade agreement status are different between pairs inside and outside a currency union (we perform the same estimator in Section 5).

Here, we address selection on an unobservable factor. Consistent with the idea that currency unions are more likely formed between countries that trade intensively, we assume that both high bilateral import shares and selection into a currency union are driven by an underlying positive shock. Vice versa, a negative shock can drive both a low bilateral import share and selection out of a currency union.

Specifically, to generate our endogenous currency union variable $\widetilde{CU}_{ij}$, we randomly choose observations of the CU_{ij} variable as observed in the data and then flip the status of non-currency union pairs to positive in response to a positive shock, and vice versa for a negative shock. We keep the mean value of the endogenous $\widetilde{CU}_{ij}$ variable the same as for CU_{ij}, and about 95% of the pairs in a currency union preserve their status.[40]

We then run a simulation to trace out the impact of currency union endogeneity. We refer to Online Appendix B where we outline our simulation procedure in more detail. In brief, we construct bilateral trade costs on the basis of trade cost function (B.1) specified in that appendix where we replace CU_{ij} with the endogenous $\widetilde{CU}_{ij}$. We then generate the simulated import shares. But crucially, we use the *same* error term for the import shares as for $\widetilde{CU}_{ij}$ to generate endogeneity between the import shares and the currency union dummy. We assume that the translog gravity model is the data generating process so that we have heterogeneous currency union effects. We run the first and second-step regressions (5) and (6) with PPML as described in Subsection 2.1.2, iterating the procedure a hundred times with fresh error terms for a sample over the period from 1990 to 2013.

Econometrically, this approach generates a positive endogeneity bias for the currency union coefficients since the bilateral trade shock is by construction correlated with the $\widetilde{CU}_{ij}$ variable. It follows that the ξ_1 main coefficient and the ξ_2 interaction coefficient in regression (6) are biased. For the ξ_1 main coefficient, we obtain a highly significant point estimate of -0.309, and for the ξ_2 interaction coefficient, we obtain a highly significant coefficient of -0.185. Both coefficients are pushed upwards.[41] The resulting currency union estimates at the mean, the 10th and the 90th percentiles follow as 0.431, 0.557 and 0.266 (all significant at the 1% level).

Figures 1 and 2 illustrate the effect of the positive endogeneity bias (we again refer to Online Appendix B for details). Both figures show the true currency union estimates in grey (those are the same across the two figures). The black lines show the estimated values (with 95% confidence intervals as dashed lines). Figure 1 plots the *unbiased* currency union estimates in the absence of endogeneity. Figure 2 plots the corresponding *biased* estimates when endogeneity is present. The estimates have roughly the same profile in both figures, indicating relatively

[40] The variables $\widetilde{CU}_{ij}$ and CU_{ij} have a correlation of around 97%.

[41] The simulated results with endogeneity are directly comparable to those without endogeneity in Table B1 where we verify our two-step procedure. In particular, the ξ_1 main coefficient of -0.309 can be compared to the -0.498 coefficient in column (3) of Table B1. The ξ_2 interaction coefficient of -0.185 can be compared to the -0.221 coefficient in column (3) of Table B1.

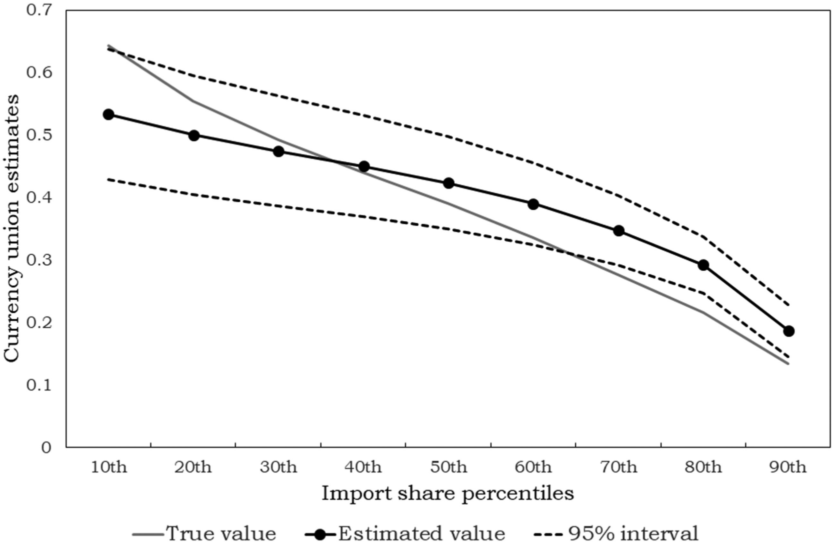

Fig. 1. *Simulated Currency Union Estimates.*
Notes: A comparison of true values (in grey) and estimated values (in black, with 95% confidence intervals as dashed lines) of currency union estimates, based on a Monte Carlo simulation for a sample over the period from 1990 to 2013. The values are reported by deciles of import shares, with the first decile denoting the lowest import shares. For example, the estimated value at the last decile (i.e., the 90th percentile) is equal to 0.187. This would imply that a currency union in the last decile is associated with an increase in bilateral trade of 21% (exp(0.187)–1 = 0.206). See Online Appendix B for details.

strong currency union effects at low import share percentiles and relatively weak effects at high percentiles. This is consistent with the empirical results in Section 2. Our estimation procedure thus correctly identifies the downward sloping profile of currency union effects, validating our approach. However, the estimates in Figure 2 are pushed upwards especially for high import shares, flattening the profile. This means that the gap between the actual and the true estimates grows for high import shares. Endogeneity thus leads to an overestimation of currency union effects for country pairs that trade intensively.

Overall, we of course do not observe the precise extent of currency union endogeneity in the actual data. But our simulation implies that if we were to correct for it, this would strengthen, rather than weaken, the heterogeneity patterns in our results. Currency union endogeneity would therefore work against us in the sense that it would make it harder to find evidence of heterogeneity patterns as in Section 2.

Finally, we also run a placebo simulation with endogenous currency unions, assuming that standard log-linear gravity is the data generating process (as opposed to translog gravity) so that by construction there is no currency union heterogeneity. We obtain currency union estimates at the mean, the 10th and the 90th percentiles of 0.298, 0.324 and 0.273 (all significant at the 1% level).[42] These estimates are biased upwards, but they do not exhibit any quantitatively meaningful pattern of currency union heterogeneity.

[42] This is in analogy to Online Appendix B.3. The currency union estimates are directly comparable to those without endogeneity in column (3) of Table B2.

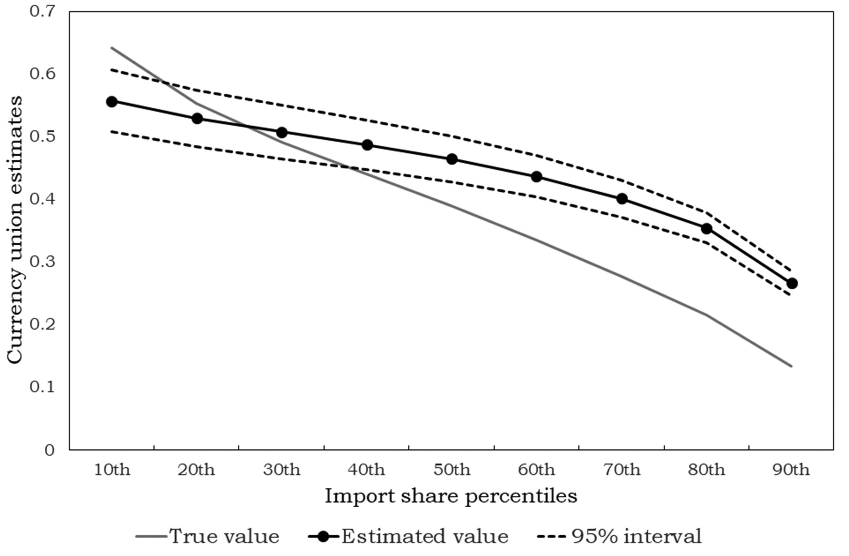

Fig. 2. *Simulated Currency Union Estimates with Endogeneity Bias.*
Notes: A comparison of true values (in grey) and estimated values (in black, with 95% confidence intervals as dashed lines) of currency union estimates subject to positive endogeneity bias, based on a Monte Carlo simulation for a sample over the period from 1990 to 2013. The values are reported by deciles of import shares, with the first decile denoting the lowest import shares. For example, the estimated value at the last decile (i.e., the 90th percentile) is equal to 0.266. This would imply that a currency union at the last decile is associated with an increase in bilateral trade of 30% (exp(0.266)–1 = 0.305). See Online Appendix B for details.

5. Robustness

To ensure the robustness of our findings, in Online Appendix C we report a battery of sensitivity checks based on alternative specifications and data samples. As in Section 2, we focus on currency unions as the main illustration. While the magnitude of the trade effect of currency unions may vary across specifications, we continue to find that it falls with bilateral import shares.

In Table C1 we apply the nearest matching estimator of Persson (2001). Our results are robust to non-random selection on observables. In Table C2 we show that currency unions continue to be associated with a heterogeneous trade effect once we control for wars, decolonisation episodes and missing data (Campbell, 2013; Campbell and Chentsov, 2020). In Table C3 we distinguish *multilateral* currency unions (i.e., between countries of similar size and wealth) from *bilateral* currency unions (i.e., when a small or poor country adopts the currency of a larger and richer country; De Sousa, 2012). Our results hold for both types of currency union.

In Table C4 we classify our currency union observations into three groups: *entry* (i.e., currency unions created during our sample period), *exit* (i.e., unions that were dissolved) and *continuous* (i.e., they existed over the whole sample period). The currency union interaction terms with log predicted shares are negative for the continuous and exit unions only. But among the entry currency unions, the common currency interaction term is negative for the euro.

Table C5 addresses the measurement of import shares per good. First, we use alternative proxies for the extensive margin n_i. We use the Hummels and Klenow (2005) measure, and we assume that the extensive margin is unity for all exporters. Second, instead of using the importing country's GDP to compute the import shares per good, we experiment with total or manufacturing gross output from the OECD STAN database (in which case our sample is reduced to nineteen OECD importing countries).

In Table C6 we consider three alternative specifications for the first-step regression (5). First, in addition to bilateral distance and contiguity we include indicator variables for sharing a common language, a common coloniser post-1945, pairs in a colonial relationship post-1945 and for territories that were, or are, part of the same country. Second, we replace bilateral distance and contiguity with (directional) country pair fixed effects. Third, we let the distance and contiguity elasticities vary over time (by interacting the two variables with year dummy variables). For the second-step regression (6), we show that our results remain robust to including time-varying distance and contiguity variables, a lagged dependent variable and a trend for EU countries or for all countries in a currency union in our sample.

In Table C7 we use alternative data samples. We use the 1949–2006 exports and GDP data from Head *et al.* (2010), and exports from the International Monetary Fund's Direction of Trade Statistics combined with GDP data from the World Bank's World Development Indicators between 1960 and 2013. We use a balanced sample between 1994 and 2013. We drop the countries (mostly island nations) omitted from the analysis of Glick and Rose (2016), the smaller nations with GDP below 500 million US dollars in 2013, the poorer countries with GDP per capita below 500 US dollars in 2013 and the post-Soviet states.

Finally, in Table C8 we test for 'feedback effects' of currency unions as discussed in Baier and Bergstrand (2007). As in their paper, we restrict our sample to five-year intervals from 1953 to 2013. Based on (6) we add one lead (i.e., values five years ahead) of the currency union dummy variable and its interaction with log predicted shares, and their coefficients are insignificant. Similar to Baier and Bergstrand (2007) we therefore do not find evidence of feedback effects.

6. Concluding Remarks

This paper offers a new approach to estimating a flexible gravity equation. Our framework has variable trade cost elasticities at its core, implying that trade costs do not always have the same trade effect across all country pairs. To introduce this form of heterogeneity, we develop a gravity framework motivated by a translog gravity equation. This approach generates variable trade cost elasticities across and within country pairs. The key prediction is that the impact of trade costs should be larger for country pairs associated with smaller import shares.

We test this prediction by employing an extensive data set of aggregate bilateral import shares for 199 countries between 1949 and 2013. We apply it to the effect of currency unions on international trade as well as a host of other trade cost variables that are popular in the literature such as the formation of regional trade agreements and membership of the WTO. Our results lend strong support to our theoretical prediction.

For example, we present evidence that the euro has promoted bilateral trade among Eurozone members but this effect is heterogeneous across country pairs. Pairs that do not trade intensively with each other tend to increase their bilateral trade by more in response to joining the euro currency union. Pairs that already have a strong trading relationship do not increase bilateral trade at all.

Regarding the trade effect of currency unions in particular, our results shed light on some of the disparities between estimates reported in the literature. By emphasising that country pairs with higher import intensity have smaller trade cost elasticities, our framework helps to explain why PPML currency union estimates are typically smaller than their OLS counterparts. In addition, as the average import share is significantly larger among euro member pairs compared to non-euro currency union pairs, our framework also helps to explain why the euro trade effect typically proves weaker on average. Our results suggest that relying on a single currency union estimate can be misleading if the objective is to assess, or to predict, the impact of a common currency on bilateral trade.

Although in our empirical application we study the trade effect of currency unions in most detail, we also demonstrate that the predictions of our theoretical framework apply more generally to a broader set of trade cost-related variables including RTAs, WTO membership, distance, sharing a common border, a common language and tariffs. These results provide strong evidence that the aggregate trade cost elasticity is variable and heterogeneous across country pairs. One potential implication is that the gains from trade liberalisation could be mismeasured if researchers erroneously assume a constant trade cost elasticity (Arkolakis *et al.*, 2012; Melitz and Redding, 2015; Bas *et al.*, 2017). Although exploring this aspect remains outside the scope of this paper, understanding the welfare implications of our results would be an important next step.

University of Warwick, UK, CAGE, UK, CESifo, Germany & CEPR, UK
University of Warwick, UK, CAGE, UK, CEP/LSE, UK, CESifo, Germany & CEPR, UK

Additional Supporting Information may be found in the online version of this article:

Online Appendix
Replication Package

References

Alesina, A. and Barro, R.J. (2002). 'Currency unions', *Quarterly Journal of Economics*, vol. 117(2), pp. 409–36.

Alesina, A., Barro, R.J. and Tenreyro, S. (2002). 'Optimal currency areas', *NBER Macroeconomics Annual*, vol. 17, pp. 301–45.

Anderson, J. and van Wincoop, E. (2003). 'Gravity with gravitas: a solution to the border puzzle', *American Economic Review*, vol. 93(1), pp. 170–92.

Arkolakis, C., Costinot, A. and Rodríguez-Clare, A. (2012). 'New trade models, same old gains?', *American Economic Review*, vol. 102(1), pp. 94–130.

Atkeson, A. and Burstein, A. (2008). 'Trade costs, pricing-to-market, and international relative prices', *American Economic Review*, vol. 98(5), pp. 1998–2031.

Baier, S.L. and Bergstrand, J.H. (2007). 'Do free trade agreements actually increase members' international trade?', *Journal of International Economics*, vol. 71(1), pp. 72–95.

Baier, S.L., Yotov, Y.V. and Zylkin, T. (2019). 'On the widely differing effects of free trade agreements: lessons from twenty years of trade integration', *Journal of International Economics*, vol. 116, pp. 206–26.

Baldwin, R. (2006). 'In or out: does it matter? An evidence-based analysis of the euro's trade effect', Centre for Economic Policy Research, London.

Baldwin, R., Di Nino, V., Fontagné, L., De Santis, R.A. and Taglioni, D. (2008). 'Study on the impact of the euro on trade and foreign direct investment', European Economy Economic Paper 321, European Commission.

Baldwin, R. and Taglioni, D. (2007). 'Trade effects of the euro: a comparison of estimators', *Journal of Economic Integration*, vol. 22(4), pp. 780–818.

Barro, R. and Tenreyro, S. (2007). 'Economic effects of currency unions', *Economic Inquiry*, vol. 45(1), pp. 1–23.

Bas, M., Mayer, M. and Thoenig, M. (2017). 'From micro to macro: demand, supply, and heterogeneity in the trade elasticity', *Journal of International Economics*, vol. 108, pp. 1–19.

Berger, H. and Nitsch, V. (2008). 'Zooming out: the trade effect of the euro in historical perspective', *Journal of International Money and Finance*, vol. 27(8), pp. 1244–60.

Berthou, A. and Fontagné, L. (2008). 'The euro effects on the firm and product-level trade margins: evidence from France', Working Paper 2008–21, CEPII.

Broda, C. and Romalis, J. (2010). 'Identifying the relationship between trade and exchange rate volatility', in (I. Takatoshi and A.K. Rose, eds.), *Commodity Prices and Markets, East Asia Seminar on Economics*, pp. 79–110, Chicago: University of Chicago Press.

Bun, M.J.G. and Klaassen, F.J.M. (2007). 'The euro effect on trade is not as large as commonly thought', *Oxford Bulletin of Economics and Statistics*, vol. 69(4), pp. 473–96.

Campbell, D.L. (2013). 'Estimating the impact of currency unions on trade: solving the Glick and Rose puzzle', *The World Economy*, vol. 36(10), pp. 1278–93.

Campbell, D.L. and Chentsov, A. (2020). 'Breaking badly: the currency union effect on trade', Working Paper 265, New Economic School.

Carrère, C., Mrázová, M. and Neary, J.P. (2020). 'Gravity without apology: the science of elasticities, distance, and trade', Economic Journal, vol. 130(628), pp. 880–910.

Chaney, T. (2008). 'Distorted gravity: the intensive and extensive margins of international trade', *American Economic Review*, vol. 98(4), pp. 1707–21.

Chen, N. and Novy, D. (2018). 'Currency unions, trade, and heterogeneity', Discussion Paper 12954, Centre for Economic Policy Research.

Correia, S., Guimarães, P. and Zylkin, T. (2020). 'Fast Poisson estimation with high-dimensional fixed effects', *Stata Journal*, vol. 20(1), pp. 95–115.

De Nardis, S. and Vicarelli, C. (2003). 'Currency unions and trade: the special case of EMU', *Review of World Economics*, vol. 139(4), pp. 625–49.

De Sousa, J. (2012). 'The currency union effect on trade is decreasing over time', *Economics Letters*, vol. 117(3), pp. 917–20.

Devereux, M.B. and Lane, P.R. (2003). 'Understanding bilateral exchange rate volatility', *Journal of International Economics*, vol. 60(1), pp. 109–32.

Eaton, J. and Kortum, S. (2002). 'Technology, geography, and trade', *Econometrica*, vol. 70(5), pp. 1741–79.

Eaton, J., Kortum, S. and Sotelo, S. (2013). 'International trade: linking micro to macro', in (D. Acemoglu, M. Arellano and E. Dekel, eds.), *Advances in Economics and Econometrics Tenth World Congress*, vol. 2, pp. 329–70, Cambridge: Cambridge University Press.

Eicher, T.S. and Henn, C. (2011). 'One money, one market: a revised benchmark', *Review of International Economics*, vol. 19(3), pp. 419–35.

Feenstra, R.C. (2003). 'A homothetic utility function for monopolistic competition models, without constant price elasticity', *Economics Letters*, vol. 78(1), pp. 79–86.

Felbermayr, G., Larch, M., Yalcin, E. and Yotov, Y.V. (2020). 'On the heterogeneous trade and welfare effects of GATT/WTO membership', Working Paper 8555, CESifo.

Flam, H. and Nordström, H. (2007). 'Explaining large euro effects on trade: the extensive margin and vertical specialization', Working Paper, Institute for International Economic Studies, Stockholm University.

Frankel, J. (2010). 'The estimated trade effects of the euro: why are they below those from historical monetary unions among smaller countries?', in (A. Alesina and F. Giavazzi, eds.), *Europe and the Euro*, pp. 169–212, Chicago: University of Chicago Press.

Frankel, J. and Rose, A. (2002). 'An estimate of the effect of common currencies on trade and income', *Quarterly Journal of Economics*, vol. 117(2), pp. 437–66.

Glick, R. (2017). 'Currency unions and regional trade agreements: EMU and EU effects on trade', *Comparative Economic Studies*, vol. 59(2), pp. 194–209.

Glick, R. and Rose, A.K. (2002). 'Does a currency union affect trade? The time-series evidence', *European Economic Review*, vol. 46(6), pp. 1125–51.

Glick, R. and Rose, A.K. (2016). 'Currency unions and trade: a post-EMU reassessment', *European Economic Review*, vol. 87, pp. 78–91.

Head, K. and Mayer, T. (2014). 'Gravity equations: workhorse, toolkit, and cookbook', in (G. Gopinath, E. Helpman and K. Rogoff, eds.), *Handbook of International Economics*, Chapter 3, pp. 131–95, Amsterdam: Elsevier.

Head, K., Mayer, T. and Ries, J. (2010). 'The erosion of colonial linkages after independence', *Journal of International Economics*, vol. 81(1), pp. 1–14.

Helpman, E., Melitz, M. and Rubinstein, Y. (2008). 'Estimating trade flows: trading partners and trading volumes', *Quarterly Journal of Economics*, vol. 123(2), pp. 441–87.

Hummels, D. and Klenow, P.J. (2005). 'The variety and quality of a nation's exports', *American Economic Review*, vol. 95(3), pp. 704–23.

Irarrazabal, A., Moxnes, A. and Opromolla, L.D. (2015). 'The tip of the iceberg: a quantitative framework for estimating trade costs', *Review of Economics and Statistics*, vol. 97(4), pp. 777–92.

Klein, M.W. (2005). 'Dollarization and trade', *Journal of International Money and Finance*, vol. 24(6), pp. 935–43.

Klein, M.W. and Shambaugh, J.C. (2006). 'Fixed exchange rates and trade', *Journal of International Economics*, vol. 70(2), pp. 359–83.

Larch, M., Wanner, J., Yotov, Y. and Zylkin, T. (2019). 'Currency unions and trade: a PPML re-assessment with high-dimensional fixed effects', *Oxford Bulletin of Economics and Statistics*, vol. 81(3), pp. 487–510.

López-Córdova, J.E. and Meissner, C.M. (2003). 'Exchange-rate regimes and international trade: evidence from the classical gold standard era', *American Economic Review*, vol. 93(1), pp. 344–53.

Mayer, T., Vicard, V. and Zignago, S. (2019). 'The cost of non-Europe, revisited', *Economic Policy*, vol. 34(98), pp. 147–99.

Melitz, M.J. and Ottaviano, G.I.P. (2008). 'Market size, trade, and productivity', *Review of Economic Studies*, vol. 75(1), pp. 295–316.

Melitz, M.J. and Redding, S.J. (2015). 'New trade models, new welfare implications', *American Economic Review*, vol. 105(3), pp. 1105–46.

Micco, A., Stein, E. and Ordoñez, G. (2003). 'The currency union effect on trade: early evidence from EMU', *Economic Policy*, vol. 37(18), pp. 315–56.

Mika, A. and Zymek, R. (2018). 'Friends without benefits? New EMU members and the 'euro effect' on trade', *Journal of International Money and Finance*, vol. 83, pp. 75–92.

Mundell, R.A. (1961). 'A theory of optimum currency areas', *American Economic Review*, vol. 51(4), pp. 657–65.

Nitsch, V. (2002). 'Honey, I shrunk the currency union effect on trade', *The World Economy*, vol. 25(4), pp. 457–74.

Novy, D. (2013). 'International trade without CES: estimating translog gravity', *Journal of International Economics*, vol. 89(2), pp. 271–82.

Parente, P.M.D.C. and Santos Silva, J.M.C. (2016). 'Quantile regression with clustered data', *Journal of Econometric Methods*, vol. 5(1), pp. 1–5.

Persson, T. (2001). 'Currency unions and trade: how large is the treatment effect?', *Economic Policy*, vol. 33(16), pp. 434–48.

Rose, A.K. (2000). 'One money, one market: estimating the effect of common currencies on trade', *Economic Policy*, vol. 30(15), pp. 8–45.

Rose, A.K. (2005). 'Which international institutions promote international trade?', *Review of International Economics*, vol. 13(4), pp. 682–98.

Rose, A.K. (2006). 'Currency unions', in (S.N. Durlauf and L.E. Blume, eds.), *The New Palgrave Dictionary of Economics*, 2nd edn, UK: Palgrave Macmillan.

Rose, A. and van Wincoop, E. (2001). 'National money as a barrier to international trade: the real case for currency union', *American Economic Review*, vol. 91(2), pp. 386–90.

Saia, A. (2017). 'Choosing the open sea: the cost to the UK of staying out of the euro', *Journal of International Economics*, vol. 108, pp. 82–98.

Santos Silva, J.M.C. and Tenreyro, S. (2006). 'The log of gravity', *Review of Economics and Statistics*, vol. 88(4), pp. 641–58.

Santos Silva, J.M.C. and Tenreyro, S. (2010). 'Currency unions in prospect and retrospect', *Annual Review of Economics*, vol. 2, pp. 51–74.

Spearot, A.C. (2013). 'Variable demand elasticities and tariff liberalization', *Journal of International Economics*, vol. 89(1), pp. 26–41.

Subramanian, A. and Wei, S.-J. (2007). 'The WTO promotes trade, strongly but unevenly', *Journal of International Economics*, vol. 72(1), pp. 151–75.