Currency Unions, Trade, and Heterogeneity*†

Natalie Chen
University of Warwick, CAGE, CESifo, and *CEPR*

Dennis Novy
University of Warwick, CAGE, CEP, CESifo, and *CEPR*

June 13, 2019

Abstract

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 the effect of currency unions on international trade. While we estimate that currency unions are associated with a trade increase of around 38 percent on average, we find substantial underlying heterogeneity. Consistent with the predictions of our framework, we find effects around three times as strong for country pairs associated with small import shares, and a zero effect for large import shares. Our results imply that conventional homogeneous currency union estimates do not provide helpful guidance for countries considering to join a currency union. Instead, countries need to take into account the distribution of their trade shares to assess the impact of trade costs.

JEL Classification: F14, F15, F33
Keywords: Currency Unions, Euro, Gravity, Heterogeneity, Trade Costs, Trade Elasticity, Translog.

---

*We are grateful to Daniel Baumgarten, 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 13^{th} 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. We thank the Economic and Social Research Council (ESRC grant ES/P00766X/1) and the Centre for Competitive Advantage in the Global Economy (CAGE, ESRC grant ES/L011719/1) at the University of Warwick for financial support.

†Natalie Chen, Department of Economics, University of Warwick, Coventry CV4 7AL, UK, CESifo and CEPR. E-mail: `N.A.Chen@warwick.ac.uk` (corresponding author), and Dennis Novy, Department of Economics, University of Warwick, Coventry CV4 7AL, UK, CEP/LSE, CESifo and CEPR. E-mail: `D.Novy@warwick.ac.uk`.

1 Introduction

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.

We apply this framework to the effect of currency unions on international trade. 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, 83 countries continued to do so. In addition, various countries are currently considering to form new currency unions or to join existing ones.[1] Countries may have several reasons to join a currency union. One of them is that currency unions are said to be associated with deeper economic integration. But does that also mean they are associated with more international trade?

To evaluate the trade effect of currency unions, researchers typically rely on a standard gravity equation framework, and insert a simple currency union dummy variable as a right-hand side regressor (e.g., Rose, 2000). This yields a single coefficient to assess the trade effect of currency unions. By construction this effect is homogeneous across all currency union country pairs in the sample.[2] A similar approach is typically taken also with regard to other trade cost variables, irrespective of whether they are represented by dummy variables.

In this paper, we challenge the view that currency unions have a homogeneous "one-size-fits-all" effect on bilateral trade flows. Our contribution is to argue theoretically, and to demonstrate empirically, that the trade effect of currency unions is heterogeneous across and within country pairs. As our theoretical framework, we introduce heterogeneous currency union effects by taking guidance from a translog gravity equation that predicts variable trade cost elasticities (Novy, 2013). In this framework, 'thin' bilateral trade relationships (characterized by small import shares) are more sensitive to trade cost changes in comparison to 'thick' trade relationships (characterized by large 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 reduction in trade costs induced by a currency union generates heterogeneous effects on trade flows, and we should therefore expect larger trade effects for country pairs associated with smaller import shares.

[1] Currency unions, or monetary unions, "are groups of countries that share a single money" (Rose, 2006). See our Data 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, the talks have stalled).

[2] General equilibrium effects might in principle differ across countries. We show in Appendix B.3 that they are quantitatively negligible and not systematically related to currency union heterogeneity.

From a methodological point of view, the flexible approach we propose in this paper can be applied more widely in the international trade literature, and we therefore hope it should be of interest to other researchers in the field. Although our paper focuses specifically on the trade effect of currency unions, the flexible gravity framework can be employed to investigate the heterogeneous trade effects of trade costs more broadly including transportation costs, tariffs, non-tariff barriers, or regional trade agreements, among others.

We start by laying out the flexible gravity framework and relate it to currency union effects in international trade. This provides 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. We adopt two main approaches to test whether the currency union effect on trade is heterogeneous across import shares. The first approach is based on the standard log-linear gravity specification that is commonplace in the literature. But instead of estimating a single currency union coefficient that is constant over the entire sample, we allow for heterogeneous currency union estimates.[3] The second approach is to estimate the translog gravity equation directly.

In the first approach, we initially run a standard gravity regression with the logarithmic import share as the dependent variable, but we then examine whether the trade effect of currency unions is heterogeneous across bilateral import shares. This form of coefficient heterogeneity creates a simultaneity bias problem, however, as the currency union effect varies with the values taken by the dependent variable. We address this issue by letting the currency union effect vary across *predicted* import shares instead. This can be achieved through a two-step methodology we propose. In the first step we generate the predicted shares by regressing the import shares on geography-related variables such as distance and contiguity, while in the second step we assess how the trade effect of currency unions varies across predicted import shares.

We carry out Monte Carlo simulations to verify the validity of our two-step procedure to estimate heterogeneous currency union effects. When we assume that the true data generating process is driven by variable trade cost elasticities, our simulations based on the two-step procedure yield results that are similar both qualitatively and quantitatively to the underlying true model. In contrast, we demonstrate that if standard gravity were the true underlying model, we would not be able to explain heterogeneous currency union effects.

When we estimate a standard gravity regression *without* heterogeneous effects, we find that sharing a common currency is associated with 38 percent more trade on average. Our contribution through the flexible gravity framework is to demonstrate that this average hides a significant amount of heterogeneity *across country pairs*. For instance, at the 90^{th} percentile of predicted import shares where shares are large, we find that the trade effect of sharing the same currency is relatively modest at 30 percent. In contrast, at the 10^{th} percentile where shares

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

are small, we find a substantially stronger effect of 94 percent. Examples of country pairs with small import shares associated with large currency union effects are Equatorial Guinea importing from Niger (105 percent), and Mali importing from the Central African Republic (98 percent). In contrast, country pairs with large import shares that do not increase trade shares at all by joining a currency union (i.e., the currency union effects are insignificant) include Belgium-Luxembourg importing from the Netherlands, Malaysia importing from Singapore, or Portugal importing from Spain.

We also find that the trade effect of currency unions is heterogeneous *within country pairs* and therefore asymmetric by direction of trade. For instance, the effect is large (at around 38 percent) when Germany imports from Portugal (i.e., low shares). But it is insignificant in the other direction when Portugal imports from Germany (i.e., large shares). As another example of within-pair asymmetry, we estimate a large currency union effect when the US imports from Panama (48 percent) but an insignificant effect when Panama imports from the US. Overall, heterogeneity in the trade effect remains robust to using different samples or specifications, controlling for non-random selection, and including the zero trade observations in the sample (Poisson Pseudo Maximum Likelihood (PPML) estimation).

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 modest.[4] Still, we find that the effect is heterogeneous across country pairs within the euro area. It is insignificant at the 90^{th} percentile of the import shares distribution, but it becomes significant and equal to 36 percent at the 10^{th} percentile. Examples of country pairs with small import shares which are associated with large euro effects are Ireland importing from Cyprus (31 percent), Finland from Malta (30 percent), and Finland from Greece (21 percent). In contrast, country pairs with large import shares not generating any additional trade from the euro include Belgium-Luxembourg importing from the Netherlands or Portugal importing from Spain.

We also explore the predictions of our model by estimating the translog gravity equation directly. In this case, the dependent variable is the bilateral import share in levels (rather than in logarithmic form). Our regressions show that sharing a common currency is associated with more bilateral trade, and consistent with our framework the magnitude of the effect falls with bilateral import shares.

One potential concern with our estimations relates to the possibly endogenous nature of currency unions. 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

[4]See Baldwin (2006), Baldwin, Di Nino, Fontagné, De Santis, and Taglioni (2008), Baldwin and Taglioni (2007), Berger and Nitsch (2008), Eicher and Henn (2011), Glick and Rose (2016), Mika and Zymek (2016), and Santos Silva and Tenreyro (2010a).

of common currencies.[5] Attempts in the literature to instrument the currency union dummy prove disappointing, however, as the instrumentation tends to increase, rather than decrease, the magnitude of currency union estimates (Alesina, Barro, and Tenreyro, 2002; Barro and Tenreyro, 2007; Rose, 2000). 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 high import intensity (i.e., large import shares) and pushing down the strong effects associated with low import intensity (i.e., small import shares). Thus, removing this potential bias would lead to even stronger heterogeneity patterns. Endogeneity can therefore not overturn our results on heterogeneity.

As they improve our understanding of how currency unions shape trade flows between trading partners, our results have important policy implications. First, by confirming that currency unions are associated with more trade between their members, our results lend support to the view that, by representing a permanent commitment to a fixed exchange rate, currency unions go beyond the simple elimination of exchange rate volatility and are likely to change the perceptions and expectations of economic agents. Second and most importantly, our results help to evaluate the potential changes in trade flows that countries can expect when joining a currency union. For instance, suppose Bulgaria, Croatia, the Czech Republic, Hungary, Poland, Romania, and Sweden were to join the euro in the next few years. As these countries are relatively small compared to some existing members of the Eurozone such as France and Germany, they have relatively large import shares. Our results suggest that these import shares will grow modestly (consistent with Baldwin, 2006; Glick, 2016; Mika and Zymek, 2016). However, trade shares in the opposite direction are smaller and can therefore be expected to grow faster.

Our paper is related to two strands of the literature. The first, and earlier one, is concerned with the trade impact of exchange rate fluctuations (for a review, see Auboin and Ruta, 2013). The received view is that by creating uncertainty, exchange rate volatility discourages trade flows. The empirical evidence is mixed, however, as the impact of volatility on trade tends to

[5]For instance, exporting and importing 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 for instance Broda and Romalis (2010) and Devereux and Lane (2003). From a theoretical point of view, see Mundell (1961) who suggests that by reducing real exchange rate fluctuations, trade reduces the costs of forming a currency union. Alesina and Barro (2002) show that countries that trade more with each other are more likely to form currency unions.

be small and not robust (Frankel and Wei, 1993).[6]

The second strand explores more specifically whether by eliminating exchange rate uncertainty, currency unions promote trade flows. In his seminal work, Rose (2000) shows that sharing a common currency more than triples bilateral trade flows. The magnitude of the effect is surprising, especially given the modest impact of exchange rate volatility estimated in the earlier literature. Subsequent work by Rose and co-authors shows that the currency union effect is smaller than initially found but remains large and robust to using different samples, specifications, and to controlling for reverse causality (Frankel and Rose, 2002; Glick and Rose, 2002; Rose and van Wincoop, 2001).[7]

These findings have inspired a large and growing literature. Various authors argue that the original "Rose" effect may be plagued by omitted variables, econometric errors, self-selection, and the presence of currency unions formed by very small or poor countries, and that the trade effect of currency unions is likely to be small or even insignificant (Baldwin, 2006; Baldwin, Di Nino, Fontagné, De Santis, and Taglioni, 2008).[8] Glick and Rose (2016) conclude that the empirical literature based on the standard gravity approach fails to deliver "reliable and robust" estimates of currency union effects since the results turn out to be highly sensitive to the econometric methodology, specification, and data sample used.[9]

Nevertheless, there is empirical evidence to suggest that heterogeneity in the trade impact of currency unions may exist along several dimensions.[10] For instance, the effect is found to be larger for developing economies (Santos Silva and Tenreyro, 2010a), smaller countries (Baldwin, 2006; Micco, Stein, and Ordoñez, 2003), and to fall over time (De Sousa, 2012). The effect is also shown to vary across currency unions (Eicher and Henn, 2011; Klein, 2005; Nitsch, 2002).[11] Consensus estimates for the euro are for instance substantially more modest than those for broader samples, falling between five and 15 percent (Baldwin, 2006; Baldwin et al., 2008). The trade effect appears stronger for industries producing highly differentiated goods

[6] The weak response of trade to exchange rate volatility could be due, among other factors, to firms hedging against exchange rate risk or the offsetting impact arising from imported inputs.

[7] Rose and Stanley (2005) carry out a meta-analysis based on the point estimates of 34 different studies and conclude that currency unions have a positive and robust effect on trade.

[8] Also see Berger and Nitsch (2008), Broda and Romalis (2010), Bun and Klaassen (2007), Campbell (2013), De Nardis and Vicarelli (2003), De Sousa (2012), Eicher and Henn (2011), Flam and Nordström (2003, 2007), Frankel (2010), Glick (2016), Klein (2005), Klein and Shambaugh (2006), Larch, Wanner, Yotov, and Zylkin (2017), López-Córdova and Meissner (2003), Micco, Stein, and Ordoñez (2003), Mika and Zymek (2016), Nitsch (2002), Persson (2001), Saia (2017), and Santos Silva and Tenreyro (2010a), among others.

[9] Baldwin et al. (2008) go as far as claiming that the empirical literature on the trade effect of currency unions "is a disaster" as the estimates reported by prominent researchers range from zero percent (e.g., Berger and Nitsch, 2008) to 1,387 percent (Alesina, Barro, and Tenreyro, 2002), most of them being "fatally flawed by misspecification and/or econometric errors."

[10] For evidence on the heterogeneous trade effects of free trade agreements, see Baier, Yotov, and Zylkin (2016) and Glick (2016). Spearot (2013) and Subramanian and Wei (2007) investigate the heterogeneous trade effects of tariff liberalization and of WTO membership, respectively.

[11] Klein (2005) finds that dollarization does not promote trade for all the countries that adopt the US dollar.

(Flam and Nordström, 2003, 2007), as well as for larger and more productive firms that adjust both at the intensive and extensive margins (Berthou and Fontagné, 2008).[12] In contrast to these papers where the various sources of heterogeneity are explored without theoretical motivation and often across different samples, we are guided by the translog gravity framework with flexible trade elasticities to derive our empirical specifications.

Our paper is organized as follows. In Section 2 we build on the translog gravity framework to motivate why we might find heterogeneous trade effects of currency unions in the data. In Section 3 we present our empirical methodology, and we describe our main results and the trade effect of the euro. We also present the translog specifications. In Section 4 we carry out Monte Carlo simulations that explore the endogeneity of currency unions. In Section 5 we provide an extensive battery of robustness checks. We conclude in Section 6. Appendix A summarizes our data and sources. In Appendix B we carry out Monte Carlo simulations that scrutinize our estimation strategy in more detail.

2 Theoretical Motivation

We use the translog gravity framework as the theoretical motivation for our analysis. As the crucial guiding feature for our purposes, this framework exhibits variable trade cost elasticities. 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 parameterization 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 $t_{ij} \neq t_{ji}$.

Imposing market clearing and solving for general equilibrium results in the translog gravity equation:

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

where x_{ij} is the bilateral trade flow between the exporting country i and the 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 thus the 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. D_i and D_j denote exporter and importer-specific terms given by:

$$D_i = \frac{y_i/y^W}{n_i} + \theta \sum_{s=1}^{S} \frac{y_s}{y^W} \ln\left(\frac{t_{is}}{T_s}\right), \tag{2}$$

$$D_j = \theta \ln(T_j), \tag{3}$$

[12]For more evidence on intensive and extensive margins, see Baldwin and Di Nino (2006), Baldwin et al. (2008), Bergin and Lin (2012), Flam and Nordström (2007), and Machado, Santos Silva, and Wei (2016).

where y^W denotes world income and S is the number of countries in the world. T_j is akin to a multilateral resistance term since it represents a weighted average of bilateral trade costs. It is given by:

$$\ln(T_j) = \sum_{s=1}^{S} \frac{n_s}{N} \ln(t_{sj}), \tag{4}$$

where N is the number of products in the world with $N \geq S$.

The translog gravity equation (1) differs in two key respects from standard gravity equations as in Anderson and van Wincoop (2003) and Eaton and Kortum (2002).[13] That is, the dependent variable is the import share per good, which means an empirical measure of n_i is required. In addition, the dependent variable is in levels, 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 examine in the context of currency unions.

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 deal with those general equilibrium effects in detail in Appendix B.3). In standard gravity equations this elasticity would be constant.[14] In the translog gravity model, however, this elasticity is variable. It follows from equation (1) as:

$$\eta_{ij} = -\frac{\theta}{\frac{x_{ij}/y_j}{n_i}}. \tag{5}$$

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 a dummy variable for currency union membership, CU_{ij}, which takes on a 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). Furthermore, we add common explanatory trade cost elements such as distance. The specified trade cost function is thus bilaterally symmetric. We provide more details in Section 3. Based on the expression in equation (5), the effect of currency union membership on trade follows as:

$$\frac{\Delta \ln(x_{ij})}{\Delta CU_{ij}} = \frac{\Delta \ln\left(\frac{x_{ij}/y_j}{n_i}\right)}{\Delta CU_{ij}} \approx -\frac{\theta\kappa}{\frac{x_{ij}/y_j}{n_i}}, \tag{6}$$

[13] See Head and Mayer (2014) for an overview of gravity equations. See Novy (2013) for a more detailed discussion of the translog gravity model.

[14] 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.

where ΔCU_{ij} indicates a switch in currency union status. Given that κ is generally negative, we expect a positive currency union effect on bilateral trade.

We would like to highlight two features of our framework. First, as the denominators of expressions (5) and (6) 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.[15] However, as we show below in our empirical analysis, we find heterogeneous effects across different pairs *within* a given currency union, and even across a given bilateral country pair by direction of trade (due to bilaterally asymmetric import shares). Second, while trade cost effects in the translog gravity framework can be bilaterally asymmetric, trade is always balanced at the aggregate country level due to the general equilibrium nature of the model.[16]

In summary, the most important insight from this motivating framework is the variable currency union effect in expression (6). More specifically, the currency union effect should be larger for country pairs associated with smaller import shares. It also follows that a symmetric reduction in bilateral trade costs induced by a currency union can lead to asymmetric increases in bilateral trade flows by direction of trade. These are testable predictions that we will now examine. While we also estimate the translog specification in equation (1) directly, we will first turn towards an alternative approach.

3 Empirical Analysis

The aim of the empirical analysis is to find out whether international trade data are characterized by variable currency union elasticities. As a starting point, we first estimate standard gravity regressions with a constant currency union elasticity so that we can compare the coefficients in our sample to those in the literature. We then proceed by exploring variable currency union elasticities. For that purpose, we adopt two approaches that are consistent with the theoretical framework in Section 2. The first approach is related to the standard log-linear gravity specification that is commonplace in the literature. But instead of estimating a single currency union coefficient that is constant over the entire sample, the novelty is that we allow for currency union effects that are heterogeneous across import shares. We explain this estimating strategy in more detail below. The second approach is to estimate the translog gravity equation (1) directly.

We use a very large, comprehensive data set of aggregate annual bilateral trade flows, covering most of global trade in modern times. It consists of an unbalanced panel including

[15] Instead of the constant κ we would then have to specify currency union-specific trade cost parameters. We allow for such an approach in our analysis of the euro in Section 3.1.4.

[16] As usual in a gravity framework, the general equilibrium adjustment in response to trade cost changes (for instance, as triggered by a new currency union) is absorbed by time-varying exporter and importer fixed effects which capture endogenous variables such as nominal income among other factors.

199 countries between 1949 and 2013. We provide details and descriptive statistics in Appendix A.

3.1 Gravity with Heterogeneity

This section describes our first approach that is related to standard log-linear gravity estimation. We initially estimate homogeneous currency union effects and then turn towards heterogeneity. We also present the trade effect of the euro. We then estimate the translog gravity equation directly.

3.1.1 Homogeneous Currency Union Effects in Standard Gravity

The dependent variable in the translog gravity equation (1) is the import share per good in levels. However, in order to obtain coefficients that we can directly compare to the literature, we first run regressions based on the standard log-linear gravity framework where the dependent variable is in logarithms. Specifically, we first estimate:

$$\ln\left(\frac{x_{ij,t}/y_{j,t}}{n_{i,t}}\right) = \alpha_1 CU_{ij,t} + \alpha_2 Z_{ij,t} + D_{i,t} + D_{j,t} + D_{ij} + \varrho_{ij,t}, \tag{7}$$

where we add time subscripts such that $x_{ij,t}$ is the FOB bilateral export value from exporter i to importer j in year t, $y_{j,t}$ is country j's nominal GDP (both in current US dollars), and $n_{i,t}$ denotes the number of goods in the exporting country which can be seen as an extensive margin measure (see details below). Trade costs depend on currency union membership, $CU_{ij,t}$, which 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 time-varying country pair variables $Z_{ij,t}$ which include dummy variables equal to one if both countries in the pair belong to an RTA or are members of the IMF, OECD, and WTO in each year, and zero otherwise (Rose, 2005).

We control for an extensive set of fixed effects. We include time-varying exporter and importer fixed effects, $D_{i,t}$ and $D_{j,t}$, which control for multilateral trade resistance and other exporter and importer-specific terms such as income. We further include country pair fixed effects D_{ij} to absorb all time-invariant bilateral trade frictions in each cross-section. These pair effects may to some extent also help to control for the endogeneity of the currency union dummy if two countries that decide to join a currency union have traditionally traded a lot with each other (but they fail to do so if the two countries decide to join following a surge in trade during the sample period, see Bun and Klaassen, 2007; Micco et al., 2003). Note that the pair effects are directional as non-directional pair effects would otherwise eliminate the asymmetry in bilateral import shares within a pair.[17] The inclusion of pair fixed effects implies that identification is achieved from the time series variation of each explanatory variable

[17] Baier et al. (2016) use non-directional pair effects to estimate the within-pair asymmetric effects of FTAs.

within each country pair.[18] For our main variable of interest, this means that identification stems from changes in bilateral currency union status over time. To control for time-invariant idiosyncratic shocks correlated at the pair level which may affect both directions of trade in a similar way (De Sousa, 2012), standard errors are clustered by non-directional country pair. The coefficients to be estimated are denoted by the α's. As sharing a common currency is expected to promote trade, we expect α_1 to be positive. The error term is $\varrho_{ij,t}$.

As the dependent variable in (7) is in logarithmic form, the $n_{i,t}$ and $y_{j,t}$ terms are absorbed by the exporter and importer fixed effects $D_{i,t}$ and $D_{j,t}$ so that in effect, we yield the same coefficients of interest as in the standard gravity specification that simply has the logarithmic bilateral trade flow on the left-hand side. However, as will become clearer later, for comparability we retain the dependent variable as specified in equation (7).

Another implication of the log-linear specification is that the zero import shares per good are excluded from the regression, and our analysis focuses on the intensive bilateral margin of adjustment. Later we show that our results remain robust to including the zero import shares in the sample by running PPML estimations (see Section 5).

To measure the exporting countries' extensive margin $n_{i,t}$, we collect each country's total exports by product category from United Nations Comtrade which are available from 1962 onwards. As the HS classification was only introduced in 1988, we rely on data at the 4-digit HS-level between 1988 and 2013, and at the 4-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 it provides us with some useful information regarding the variation in the extensive margin across exporting countries.

We check the robustness of our findings by using alternative proxies for the extensive margin. First, we rely on the cross-country measure constructed by Hummels and Klenow (2005), using export data from 126 exporting to 59 importing countries in more than 5,000 6-digit HS-level product categories for 1995. Second, as these authors observe that the extensive margin tends to be stronger for larger economies, we rely on the GDP of the origin country as an alternative proxy.[19] Finally, we also report results where we assume that the extensive margin is unity for all exporters, in which case the dependent variable is simply the bilateral import share.

[18] The recent literature concludes that time-varying exporter and importer dummies and time-invariant country pair fixed effects should be included (Baldwin, 2006; Baldwin and Taglioni, 2007; Baldwin et al., 2008; De Nardis and Vicarelli, 2003; Eicher and Henn, 2011; Mika and Zymek, 2016). The earlier literature failed to do so (e.g., Rose, 2000).

[19] In that case, the proxy for the exporter's extensive margin is time varying.

3.1.2 Heterogeneous Currency Union Effects

Our ultimate aim is to investigate whether the trade effect of currency unions, as captured by α_1 in specification (7), is heterogeneous across bilateral import shares per good. If we simply allowed α_1 to vary by 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 let the currency union effect vary across *predicted* import shares. For that purpose, we adopt 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 these predicted shares. We now explain this approach in more detail.

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

$$\ln\left(\frac{x_{ij,t}/y_{j,t}}{n_{i,t}}\right) = D_{i,t} + D_{j,t} + \delta K_{ij} + \nu_{ij,t}, \tag{8}$$

where K_{ij} includes geography-related variables, i.e., logarithmic bilateral distance and a contiguity dummy. We do not include the time-varying pair variables for RTAs, currency unions, the OECD, IMF, or WTO as they are not geography-related and therefore more likely endogenous. We then generate the predicted shares, denoted by $\widehat{\ln\left(\frac{x_{ij,t}/y_{j,t}}{n_{i,t}}\right)}$.[20]

In the second step we include an interaction term between the currency union dummy and the predicted import shares, with ξ_2 as the key coefficient of interest. We estimate:

$$\ln\left(\frac{x_{ij,t}/y_{j,t}}{n_{i,t}}\right) = \xi_1 CU_{ij,t} + \xi_2 CU_{ij,t} \times \widehat{\ln\left(\frac{x_{ij,t}/y_{j,t}}{n_{i,t}}\right)} + \xi_3 Z_{ij,t} + D_{i,t} + D_{j,t} + D_{ij} + \varepsilon_{ij,t}. \tag{9}$$

Since this specification includes exporter-year, importer-year and country pair fixed effects, the main effect of the predicted import shares drops out from the regression. 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 predicted import share is a generated regressor, we bootstrap standard errors with 100 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

[20] As we show later, our results are robust to including further gravity controls in specification (8), or to simply controlling for time-invariant country pair fixed effects. Our results also remain robust to including the time-varying pair variables for RTAs, currency unions, the OECD, IMF, or WTO in the first step, which is akin to running an instrumental variables regression.

(distance and contiguity) as well as time-varying exporter and importer-specific fixed effects, and then predicting the import shares. In 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 (we show that correcting for currency union endogeneity would make the heterogeneity patterns even stronger).

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:

$$\ln\left(\frac{x_{ij,t}/y_{j,t}}{n_{i,t}}\right) = \beta_{1,h}CU_{ij,t} \times D_h + \beta_2 Z_{ij,t} + D_{i,t} + D_{j,t} + D_{ij} + D_h + \epsilon_{ij,t}, \tag{10}$$

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. The regression also includes interval fixed effects, D_h. Consistent with our theoretical framework, we expect the currency union effect to be largest in the interval of the lowest predicted shares, and to be weaker in intervals of higher shares.[21]

Recall that due to their logarithmic form, the estimation of equations (9) or (10) yields exactly the same coefficients regardless of whether we use the logarithmic import share, the logarithmic import share per good, or the logarithmic bilateral trade flow as the dependent variable. The reason is that the exporter-year and importer-year fixed effects absorb all country-specific variables such as $n_{i,t}$ and $y_{j,t}$. In contrast, note that for the first-step regression (8), we must use the import share per good as the dependent variable since this is the variable we need to predict.

3.1.3 Baseline Results

We start by discussing homogeneous currency union estimates. In column (1) of Table 1, we estimate equation (7) but only include the currency union dummy as a regressor.[22] Its estimated coefficient is equal to 0.363, suggesting that a common currency is associated with an increase in bilateral trade of 44 percent on average ($\exp(0.363) - 1 = 0.438$). When we add the time-varying country pair controls in column (2), the magnitude of the effect slightly decreases to 38 percent. Belonging to an RTA, and membership of the OECD, IMF and WTO

[21] According to the theoretical framework, the effects of the RTA, IMF, OECD, and WTO variables should also be heterogeneous across import shares. We show that our results remain robust to also interacting these variables with the predicted shares in equation (9), and to estimating their effects separately by intervals of predicted shares in equation (10). 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 standard errors. The magnitude of the currency union effect is overestimated due to the omission of the fixed effects, but we still find evidence that the effect falls with bilateral import shares.

[22] In all tables, the number of observations differs from the total number of observations in the sample because the observations that are perfectly predicted by the fixed effects (i.e., singletons) are dropped.

all have a positive association with bilateral trade (Rose, 2005). We note that these currency union estimates capture direct trade cost effects. We discuss indirect general equilibrium effects in Appendix B.3.[23]

Table 1: Baseline Results

	(1)	(2)	(3)	(4)	(5)
CU	0.363^a (0.058)	0.326^a (0.057)	-0.075 (0.109)	-0.040 (0.107)	0.066 (0.110)
CU×predicted share	–	–	-0.073^a (0.018)	-0.061^a (0.017)	-0.047^a (0.018)
RTA	–	0.415^a (0.028)	–	0.413^a (0.026)	–
IMF	–	0.165^b (0.065)	–	0.164^a (0.053)	–
OECD	–	0.366^a (0.051)	–	0.365^a (0.037)	–
WTO	–	0.146^a (0.035)	–	0.144^a (0.030)	–
CU estimates					
Mean	–	–	0.525^a (0.072)	0.462^a (0.071)	0.453^a (0.071)
10^{th} percentile	–	–	0.765^a (0.120)	0.663^a (0.117)	0.608^a (0.118)
90^{th} percentile	–	–	0.288^a (0.054)	0.263^a (0.053)	0.300^a (0.053)
R-squared	0.807	0.808	0.807	0.808	0.808
Observations	780,818	780,818	780,818	780,818	780,818

Notes: Exporter-year, importer-year, and (directional) country pair fixed effects are included. Robust standard errors adjusted for clustering at the (non-directional) country pair level are reported in parentheses in (1) and (2). Standard errors are bootstrapped in (3) to (5). In (5), the time-varying country pair controls are interacted with predicted import shares per good (not reported). ^{a} and ^{b} indicate significance at the one and five percent levels, respectively. The dependent variable is the log import share per good. "predicted share" is the predicted log import share per good.

Our next task is to demonstrate whether these results mask some heterogeneity in the trade effect of currency unions across country pairs. To do so, we first run the first-step regression (8). As expected, import intensity is stronger between closer and contiguous countries (the estimated coefficients are significant at the one percent level). We then generate the predicted import shares per good and run the second-step regression (9).

Columns (3) and (4) report the same specifications as columns (1) and (2) but include an interaction between the currency union dummy and the predicted shares. In column (5) we further interact the time-varying pair controls with the predicted shares (not reported).[24] In all cases, the interaction terms with the currency union dummy are negative and significant.

[23] As we show, those are second-order effects that are quantitatively small. The intuition is that currency unions are relatively rare at the bilateral level (see the Data 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.

[24] Also consistent with our model's predictions, the interactions between the time-varying pair controls and the predicted import shares are negative and significant (with the exception of the IMF and WTO variables).

Therefore, the impact of currency unions is heterogeneous. More specifically, it falls with bilateral import shares.[25]

In the lower part of Table 1 in columns (3) to (5), we report the implied currency union estimates at the mean and at different percentiles of the predicted import shares distribution as well as their corresponding standard errors. The magnitude goes down when we move from the 10^{th} to the 90^{th} percentile of the predicted import shares. For instance in column (4), the currency union estimate is equal to 0.462 at the mean value of predicted shares, 0.663 for a country pair at the 10^{th} percentile, and 0.263 at the 90^{th} percentile. In other words, at the 10^{th} percentile currency unions are associated with 94 percent more trade ($\exp(0.663) - 1 = 0.941$), whereas at the 90^{th} percentile the corresponding effect is only 30 percent.

Table 2: Pair Specific Currency Union Effects and Bilateral Asymmetries

Pair Specific Currency Union Effects					
Large CU Effects			Small CU Effects		
Exporter	Importer	CU estimates	Exporter	Importer	CU estimates
Niger	Eq. Guinea	0.719^a (0.132)	Netherlands	Belg.-Lux.	−0.148 (0.134)
Central Afr. Rep.	Mali	0.683^a (0.123)	Germany	Netherlands	−0.124 (0.128)
Chad	Côte d'Ivoire	0.661^a (0.117)	Singapore	Malaysia	−0.093 (0.120)
Liberia	The Bahamas	0.657^a (0.116)	Spain	Portugal	0.037 (0.088)
Bilateral Asymmetries					
Large CU Effects			Small CU Effects		
Exporter	Importer	CU estimates	Exporter	Importer	CU estimates
Panama	United States	0.394^a (0.060)	United States	Panama	0.074 (0.080)
Portugal	Germany	0.323^a (0.053)	Germany	Portugal	0.111 (0.073)
Greece	France	0.376^a (0.058)	France	Greece	0.186^a (0.060)
France	Germany	0.027 (0.091)	Germany	France	0.010 (0.095)

Notes: Currency union estimates evaluated at 2013 predicted import shares using the estimated coefficients reported in column (4) of Table 1. Bootstrapped standard errors adjusted for clustering at the (non-directional) country pair level are reported in parentheses. ^{a} indicates significance at the one percent level.

As an illustration, the upper part of Table 2 reports examples of country pairs with large and small currency union estimates (evaluated at 2013 predicted import shares). 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 (105 percent), Mali from the Central African Republic (98 percent), Côte d'Ivoire from Chad (94 percent), and the Bahamas from Liberia (93 percent). In contrast, some country pairs with large import shares do not increase trade shares by joining a currency union, the effect being insignificant

[25] Irarrazabal, Moxnes, and Opromolla (2015) introduce 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.

for Belgium-Luxembourg importing from the Netherlands, the Netherlands from Germany, Malaysia from Singapore, or Portugal from Spain.

The lower part of Table 2 shows that the trade effect of currency unions is also heterogeneous *within country pairs* and therefore asymmetric by direction of trade. For instance, the effect is large (at 48 percent) when the US imports from Panama (i.e., a low predicted share) but insignificant when Panama imports from the US (i.e., a large predicted share). The effect is also large when Germany imports from Portugal, or France imports from Greece (low predicted shares) but insignificant or small in the other direction (large predicted shares). In contrast, as France and Germany both import intensively from each other, sharing the same currency has no effect on their bilateral trade in either direction.

Let us consider the example of Germany and Portugal in more detail. For imports from Portugal to Germany, the import share is low at 0.2 percent for the year 2013. But in the opposite direction, the import share is relatively large at 4.5 percent. The corresponding trade flow values are \$7.7bn and \$9.4bn (i.e., Portugal has a bilateral trade deficit with Germany). Suppose these two countries would no longer be in a currency union, i.e., their bilateral currency union dummy would switch from 1 to 0, and bilateral trade costs would go up. Then according to the expression in equation (6) and the estimates in Table 2, the import share from Portugal to Germany would decrease by 38 percent all else being equal.[26] In the data this would imply a reduction of trade from Portugal to Germany by about \$2.9bn. The import share in the other direction would decrease by 12 percent, corresponding to a reduction of trade from Germany to Portugal by roughly \$1.1bn. Thus, the bilateral trade deficit would widen.[27] Vice versa, a reduction in bilateral trade costs would shrink the bilateral trade deficit. We note that although aggregate trade is balanced, the translog gravity framework in equation (1) is consistent with bilateral imbalances even if bilateral trade costs are symmetric as in our example.[28]

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 Appendix B we carry out Monte Carlo simulations to verify the validity of our procedure. When we assume that the true data generating process is driven by variable trade cost elasticities as suggested by our theoretical framework in Section 2, our simulations based on the two-step procedure yield results that are similar both qualitatively and quantitatively to the underlying true model. In contrast, we demonstrate that if standard log-linear gravity were the true underlying model, we would not be able to explain heterogeneous currency union effects.

[26] This is calculated as $\exp(0.323) - 1 = 0.381$.

[27] While the estimates in Table 2 are based on import shares per good, in our data the total import shares for Germany and Portugal are roughly the same multiple of import shares per good. For simplicity, we therefore frame the example in terms of total import shares.

[28] Aggregate trade remains balanced through general equilibrium adjustments. That is, apart from the direct trade cost effect mentioned in the text, indirect price index effects would also be in operation (in particular, see equations 2 and 3). Incomes would ultimately also adjust in general equilibrium.

Table 3: Heterogeneous Currency Union Effects: Intervals

	(1)	(2)	(3)	(4)	(5)
CU (first interval)	0.481^{a} (0.093)	0.630^{a} (0.137)	0.696^{a} (0.161)	0.718^{a} (0.160)	0.496^{a} (0.117)
CU (second interval)	0.285^{a} (0.059)	0.415^{a} (0.079)	0.445^{a} (0.098)	0.413^{a} (0.097)	0.428^{a} (0.087)
CU (third interval)	–	0.249^{a} (0.064)	0.347^{a} (0.076)	0.341^{a} (0.076)	0.213^{a} (0.080)
CU (fourth interval)	–	–	0.253^{a} (0.067)	0.281^{a} (0.067)	0.308^{a} (0.073)
RTA	0.415^{a} (0.028)	0.409^{a} (0.028)	0.403^{a} (0.028)	–	0.406^{a} (0.028)
IMF	0.164^{b} (0.065)	0.168^{a} (0.064)	0.169^{a} (0.064)	–	0.168^{a} (0.064)
OECD	0.366^{a} (0.051)	0.358^{a} (0.051)	0.349^{a} (0.051)	–	0.351^{a} (0.051)
WTO	0.146^{a} (0.035)	0.146^{a} (0.035)	0.146^{a} (0.035)	–	0.146^{a} (0.035)
Intervals split by	# obs.	# obs.	# obs.	# obs.	# obs. CU=1
R-squared	0.808	0.808	0.808	0.808	0.808
Observations	780,818	780,818	780,818	780,818	780,818

Notes: Exporter-year, importer-year, (directional) country pair, and interval fixed effects are included. Robust standard errors adjusted for clustering at the (non-directional) country pair level are reported in parentheses. In (4), the time-varying country pair controls are interacted with the interval fixed effects (not reported). ^{a} and ^{b} indicate significance at the one and five percent levels, respectively. The dependent variable is the log import share per good.

We proceed by regressing equation (10). In Table 3 we report currency union effects estimated separately by intervals of predicted import shares per good. Based on the median of the predicted shares distribution, column (1) of Table 3 splits the data into two distinct intervals, where the first and second intervals refer to the intervals with the lowest and highest shares, respectively. As expected, the currency union coefficient is largest (equal to 0.481) for the lowest shares and smallest (equal to 0.285) for the largest shares. Columns (2) and (3) split the sample into three and four equally-sized intervals of predicted import shares per good. In both cases the magnitude of the currency union coefficient gradually declines from the first to the last interval.[29] In column (4) we also let the effects of the time-varying pair controls vary across the four intervals (not reported). The results remain qualitatively similar. Finally, in column (5) we split the data into four intervals but in such a way that each includes roughly the same number of observations for which the currency union dummy is equal to one.[30] Compared to column (3), the currency union estimates are slightly smaller but their magnitude continues to fall with predicted shares (with the exception of the interval capturing the lowest shares).[31]

[29] In column (3) the number of observations for which the currency union dummy is equal to one is 1,379 in the first, 3,171 in the second, 3,816 in the third, and 4,719 in the fourth interval.

[30] There are 3,271 observations for which the currency union dummy is equal to one in the first three intervals and 3,272 in the fourth.

[31] We can reject that the currency union coefficients are equal across intervals at the five percent level in columns (1) and (2), and at the ten percent level in columns (3) to (5).

3.1.4 The Euro

Given the prominence of the European single currency, we investigate the trade effect of the euro in more detail. We start by providing homogeneous estimates. In column (1) of Table 4 we first run the standard specification (7) on the full sample, but the currency union dummy is split between euro and non-euro currencies. Sharing a common currency is associated with 50 percent more trade for the euro ($\exp(0.403) - 1 = 0.496$), and with 31 percent more trade for non-euro currencies ($\exp(0.270) - 1 = 0.310$). This suggests that compared to other common currencies, the positive trade impact of the euro may be larger.[32]

Table 4: The Euro

	(1)	(2)	(3)		(4)	
CU non EURO	0.270^a (0.087)	0.270^a (0.087)	-0.718^a (0.226)		-0.653^a (0.227)	
CU non EURO×predicted share	–	–	-0.144^a (0.030)		-0.135^a (0.030)	
EURO	0.403^a (0.063)	0.023 (0.069)	0.471^a (0.094)		-0.194^c (0.110)	
EURO×predicted share	–	–	0.015 (0.020)		-0.044^b (0.021)	
RTA	0.414^a (0.028)	0.400^a (0.028)	0.409^a (0.026)		0.394^a (0.026)	
IMF	0.165^b (0.064)	0.172^a (0.065)	0.167^a (0.052)		0.173^a (0.052)	
OECD	0.365^a (0.051)	0.343^a (0.051)	0.361^a (0.037)		0.338^a (0.037)	
WTO	0.147^a (0.035)	0.157^a (0.035)	0.149^a (0.030)		0.158^a (0.030)	
Trend EU countries	–	0.027^a (0.003)	–		0.028^a (0.003)	
CU estimates			Non EURO	EURO	Non EURO	EURO
Mean	–	–	0.466^a (0.084)	0.349^a (0.109)	0.453^a (0.083)	0.167 (0.112)
10^{th} percentile	–	–	0.939^a (0.154)	0.301^c (0.166)	0.896^a (0.151)	0.311^c (0.169)
90^{th} percentile	–	–	−0.003 (0.099)	0.397^a (0.067)	0.015 (0.101)	0.024 (0.074)
R-squared	0.808	0.808	0.808		0.808	
Observations	780,818	780,818	780,818		780,818	

Notes: Exporter-year, importer-year, and (directional) country pair fixed effects are included. Robust standard errors adjusted for clustering at the (non-directional) country pair level are reported in parentheses in (1) and (2). Standard errors are boostrapped in (3) and (4). ^{a}, ^{b}, and ^{c} indicate significance at the one, five, and ten percent levels, respectively. The dependent variable is the log import share per good. "predicted share" is the predicted log import share per good.

However, as argued by previous authors, one issue with the regression in column (1) is that it fails to control for the effect of European Union 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 (2) we further include a pair-specific time trend for EU countries

[32] If we run equation (7) on the full sample but excluding the Eurozone, the trade effect of sharing a common currency is 30 percent on average, compared to 38 percent in the full sample (column 2 of Table 1).

(both in and out of the euro) to control for the ongoing European integration process (Baldwin, 2006; Baldwin et al., 2008; Berger and Nitsch, 2008; Bun and Klaassen, 2007; Micco et al., 2003; Mika and Zymek, 2016). The positive coefficient on the trend indicates that on average, EU countries trade more intensively with each other over time.[33] Interestingly, the inclusion of the trend turns the euro effect insignificant.[34] Berger and Nitsch (2008), Bun and Klaassen (2007), and Mika and Zymek (2016) also find that the inclusion of a time trend dramatically reduces the magnitude and significance of the euro trade effect.[35] As the average import share per good in our sample is significantly larger for euro member pairs compared to non-euro currency union pairs, finding that the euro trade effect is weaker on average is thus consistent with the theoretical framework in Section 2.[36]

We now turn to heterogeneity. Columns (3) and (4) report the same specifications as columns (1) and (2) but we interact the currency union indicators with the predicted import shares. In both cases we observe heterogeneity in the trade effect of non-euro currency unions. For the euro, the interaction term is negative and significant in column (4) only once the pair-specific EU trend is included. Put differently, although the euro effect is on average insignificant, it is heterogeneous across country pairs, and it is larger for the pairs associated with smaller import shares.[37]

As shown in the lower part of Table 4, while in column (4) the euro estimate is insignificant at the mean value of predicted shares, it is equal to 0.311 for a country pair at the 10^{th} percentile of the predicted shares distribution. Examples of country pairs with small import shares which are associated with large trade effects of the euro are Ireland importing from Cyprus (31 percent), Finland from Malta (30 percent), and Finland from Greece (21 percent). In contrast, country pairs with large import shares which do not increase trade shares through the euro include Belgium-Luxembourg importing from the Netherlands or Portugal importing from Spain (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 Germany imports from

[33] The trend controls for EU-related policies including the Single Market, treaties on EU integration, the Exchange Rate Mechanism, etc. It is included for 27 EU countries (as Belgium and Luxembourg are merged) and for the overseas territories of the EU including French Guiana, Greenland (between 1973 and 1985 only), Guadeloupe, Martinique, and Réunion (Gibraltar belongs to the EU but is excluded from our data set). For each country pair the trend starts in the year a country joins the EU. Another way of controlling for the effect of EU integration is to run regressions for the EU15 or the EU28 only (Baldwin, 2006). Also see Saia (2017) who employs a synthetic control method to assess the trade effect of the euro.

[34] If we include a pair-specific trend for Eurozone countries only, the euro coefficient becomes smaller but remains significant. Including a trend for EU countries is, however, more appropriate as EU integration has affected all EU countries whether or not they adopted the euro (Baldwin, 2006).

[35] Our results remain similar if we exclude from the sample the small overseas territories which are associated with EU countries and use the euro (in our sample these are Saint Pierre et Miquelon, French Guiana, Guadeloupe, Martinique, and Réunion). Note that Andorra, Kosovo, Montenegro, and San Marino, neither of which belong to the EU, also use the euro. Only Andorra is included in our sample.

[36] The average import shares per good are 2 percent and 1.4 percent, respectively.

[37] We get similar results if we assume that the euro was introduced in 2002 (as a paper currency) as opposed to 1999 (as an electronic currency).

Cyprus (i.e., low predicted shares). But it is insignificant when Cyprus imports from Germany (i.e., 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 modest. Our main interest instead lies in the heterogeneous effects of currency unions. Consistent with the predictions of our model we find that the effect of the euro is heterogeneous across and within country pairs.

3.2 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 equation (6). As the dependent variable in (1) is in levels (rather than in logarithmic form), the zero trade observations can be included in the sample. We therefore report two sets of results, i.e., excluding and including the zeros. We also note that in contrast to the two-step OLS regressions reported earlier, the translog approach does not require us to predict the bilateral import shares per good in a first step.

Columns (1) and (2) of Table 5 report the results excluding and including the zero observations. In column (1) the currency union coefficient is equal to 0.006. As shown in the lower part of the table, this corresponds to an estimate of 0.912 at the mean value of import shares. Therefore, sharing a common currency is associated with 91 percent more bilateral trade, a magnitude which is larger than the average of 38 percent estimated in the standard gravity framework (column 2 of Table 1). As predicted by the translog framework, the effect is heterogeneous across country pairs, and the currency union estimates decrease from the 10^{th} to the 90^{th} percentile of import shares per good. However, we note that the currency union estimates at the 10^{th} percentile are extremely large compared to previous tables. The reason is that translog imposes a hyperbolic functional form on the way the currency union elasticities are computed. This can be seen in equation (6) 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 the Data Appendix A), the implied elasticities tend to become very large in a mechanical fashion. We therefore treat the currency union estimates at low percentiles with a particular degree of caution.

Besides, all other regressors are significant and 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 magnitude of the currency union effect at the mean value of (non-zero) import shares is smaller at 47 percent, but qualitatively our results continue to hold. Sharing a common currency is associated with more bilateral trade, and this effect is stronger for the country pairs with smaller import shares.

Table 5: Translog Estimation

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

Notes: Exporter-year, importer-year, and (directional) country pair fixed effects are included. Robust standard errors adjusted for clustering at the (non-directional) country pair level are reported in parentheses. [a] and [b] indicate significance at the one and five percent levels, respectively. The dependent variable is the import share per good in levels.

4 Endogeneity

Currency unions are not randomly assigned. Santos Silva and Tenreyro (2010a) argue 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) addresses selection on observables. By applying a propensity-score matching estimator, he accounts 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 take the CU_{ij} variable as observed in the data and combine it with a randomly drawn error term in an additive way. We round the resulting values to 0 and 1 to preserve the dummy nature of this variable. Overall, 95 percent of the pairs in a currency union preserve their status, and we keep the mean value of the endogenous $\widetilde{CU}_{ij}$ variable roughly the same as for CU_{ij}.[38]

[38] $\widetilde{CU}_{ij}$ and CU_{ij} have a correlation of around 97 percent.

We then run a simulation to trace out the impact of currency union endogeneity. We refer to 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 (B1) specified in that appendix where we replace CU_{ij} with the endogenous $\widetilde{CU}_{ij}$. We then generate the simulated import shares.[39] 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 true data generating process so that we have heterogeneous currency union effects. We run the first and second-step regressions (8) and (9) as described in Section 3.1.2, iterating the procedure 100 times with fresh error terms.

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 (9) are pushed upwards. For the ξ_1 main coefficient we obtain a highly significant point estimate of 0.235, and for the ξ_2 interaction coefficient we obtain a coefficient of -0.063, significant at the five percent level. Both coefficients are subject to upward bias.[40] The resulting currency union estimates at the mean, 10^{th} and the 90^{th} percentiles follow as 0.478, 0.573, and 0.399 (all significant at the one percent level but not significantly different from each other). Compared to our baseline results in Table 1, the heterogeneity profile is therefore flattened, with small import shares no longer being associated with currency union estimates that are statistically different from those for large import shares.

Figures 1 and 2 illustrate the effect of positive endogeneity bias for a particular simulation (we again refer to Appendix B for details). The black lines (with 95 percent confidence intervals in dashed lines) are directly comparable. Figure 1 plots the (unbiased) estimated currency union estimates in the absence of endogeneity where the estimates are strong at low import share percentiles, and weak or zero for large import shares at high percentiles. Figure 2 plots the corresponding (biased) estimates when endogeneity is present. In that case, the currency union estimates are no longer statistically different across import share percentiles. Thus, endogeneity bias works to flatten the heterogeneity profile. However, the key point to note is that in our empirical results based on actual data in Section 3, we do not find a flat currency union effect profile. This means that if we were to correct for endogeneity bias in the actual data (to the extent that it exists), this would strengthen, rather than weaken, the heterogeneity patterns in our results.

We also run placebo simulations with endogenous currency unions, assuming that standard log-linear gravity is the true data generating process (as opposed to translog gravity) so that

[39] We use the same data sample as in Section B.1.

[40] 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 0.235 can be compared to the -1.103 coefficient in column (3) of Table B1. The ξ_2 interaction coefficient of -0.063 can be compared to the -0.363 coefficient in that column.

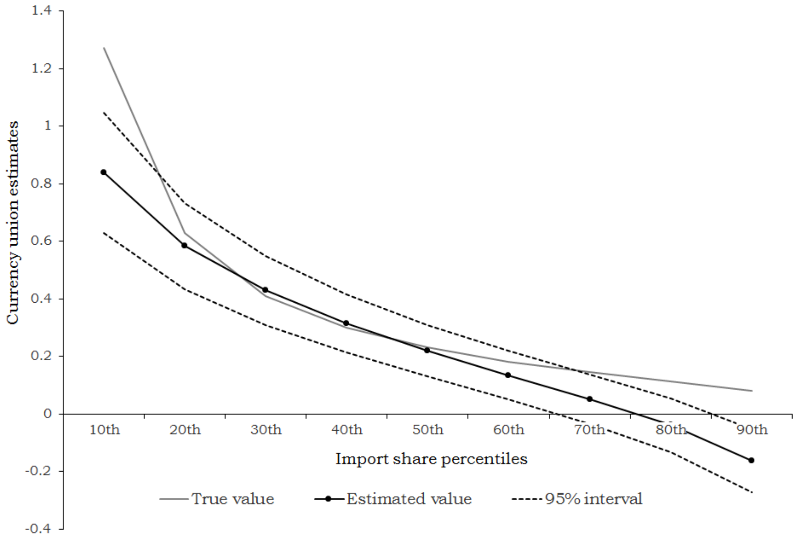

Figure 1: A comparison of true values (in grey) and estimated values (in black, with 95 percent confidence intervals as dashed lines) of currency union estimates, based on Monte Carlo simulation. The values are reported by deciles of import shares, with the first decile denoting the lowest import shares. For example, the estimated value at the first decile (i.e., 10^{th} percentile) is equal to 0.839. This would imply that a currency union in the first decile is associated with an increase in bilateral trade of 131 percent (exp(0.839)–1 = 1.31). See Appendix B for details.

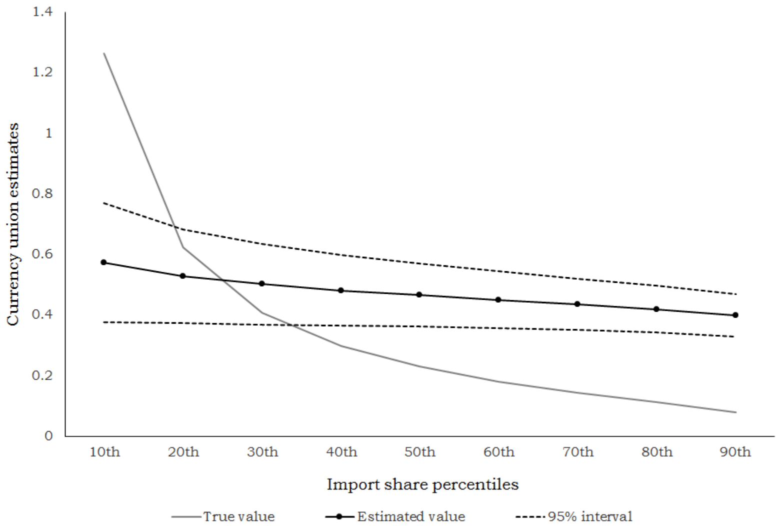

Figure 2: A comparison of true values (in grey) and estimated values (in black, with 95 percent confidence intervals as dashed lines) of currency union estimates subject to positive endogeneity bias, based on Monte Carlo simulation. Endogeneity bias flattens the heterogeneity profile compared to Figure 1. 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 first decile (i.e., 10^{th} percentile) is equal to 0.573. This would imply that a currency union at the first decile is associated with an increase in bilateral trade of 77 percent (exp(0.573)–1 = 0.77). See Appendix B for details.

there is by construction no currency union heterogeneity. We obtain currency union estimates at the mean, 10^{th} and the 90^{th} percentiles of 0.909, 0.471, and 1.351 (all significant at the one percent level and significantly different from each other).[41] Not only are these estimates biased upwards, they also exhibit a heterogeneity pattern that goes in the *opposite* direction of what our theory predicts. That is, larger import shares are associated with larger currency union effects. Clearly, this is not the pattern we find in the actual data.

In summary, we conclude that endogeneity either dampens or in some cases even overturns the heterogeneity patterns predicted by theory. The endogeneity of currency unions – to the extent that it exists – would therefore work against us and make it harder to find evidence of heterogeneity patterns as we do in Section 3.

5 Robustness

To ensure the robustness of our findings, this section discusses our results based on alternative specifications and data samples. Despite some variation in the magnitude of the trade effect of currency unions across specifications, we find heterogeneous effects across country pairs, thus supporting our paper's main conclusions. The results are reported in Appendix C.

Zero Trade Observations One way of dealing with the zero trade observations is to estimate our regressions by PPML, using the import shares per good in levels as the dependent variable (Santos Silva and Tenreyro, 2006).[42]

In column (1) of Table C1, we regress the import shares per good in levels on the controls of equation (7) but first exclude the zero observations. The currency union dummy variable indicates that sharing a common currency is associated with 22 percent more trade on average ($\exp(0.202) - 1 = 0.224$). Column (2) shows that the effect is heterogeneous across country pairs since we find a negative interaction between the currency union dummy and the predicted shares.[43] This shows that when we exclude the zero observations from the sample, the OLS and PPML estimations yield very similar results.[44]

We proceed by including the zero observations in the sample. Columns (3) and (4) report the same specifications as the two previous columns. According to column (3), sharing a common currency is associated with 29 percent more trade. Evidence of heterogeneity again arises in column (4) where we interact the currency union dummy with the predicted shares.[45] Finally,

[41] This is in analogy to Section B.2. The currency union estimates are directly comparable to those without endogeneity in column (3) of Table B2. They are far from the true values in column (1) of Table B2.

[42] We employ the `ppmlhdfe` Stata command written by Correia, Guimarães, and Zylkin (2019). It estimates a Poisson pseudo-likelihood regression allowing for multiple levels of fixed effects.

[43] The predicted shares are constructed based on a PPML specification analogous to (8).

[44] In columns (2), (4), and (5), we boostrap standard errors with 100 replications.

[45] The elasticities at the mean, 10^{th} and the 90^{th} percentiles are calculated for non-zero import shares only.

in column (5) we distinguish between the trade effects of euro and non-euro currencies. Again, our results remain largely similar, and the trade effects of both euro and non-euro currencies are heterogeneous across country pairs.[46]

In sum, these results suggest that incorporating the zero observations in the sample does not qualitatively affect our conclusions, which support the notion that the trade effect of currency unions falls with bilateral import shares. This contrasts with the results of De Sousa (2012), Glick and Rose (2016), and Mika and Zymek (2016) who find that OLS and PPML currency union estimates significantly diverge from each other.[47]

Selection on Observables Persson (2001) claims that the trade impact of common currencies can be mismeasured if the countries in a currency union are systematically different from those outside (see, also, Baldwin, 2006; Baldwin et al., 2008). To address this issue Persson (2001) applies a matching technique to identify the non-currency union country pairs that are most similar to the currency union pairs. He then compares bilateral trade flows between currency union members and their matched non-currency union counterparts. He finds that the trade effect of currency unions is insignificant. In contrast, Rose (2001) provides evidence that the magnitude of his currency union estimates remains robust to matching. To check whether non-random selection matters for our results, we apply the nearest matching estimator of Persson (2001) and Rose (2001). We run a probit regression to generate the propensity score, and match the currency union observations with the non-currency union observations that deviate by no more than a small distance from the propensity score.[48]

We estimate specifications (7) and (9) on the matched sample.[49] Column (1) of Table C2 reports the results for equation (7). On average, currency unions are associated with 42 percent more trade. The trade effect of currency unions is heterogeneous across country pairs, as shown in column (2) which interacts the currency union indicator with predicted import shares per good. We therefore conclude that our findings remain robust to non-random selection on observables.

[46] In all cases, the RTA, IMF, and OECD variables remain positive and significant while the WTO dummy is insignificant (not reported).

[47] Only Mika and Zymek (2016) report PPML estimates with country pair and time-varying exporter and importer fixed effects. De Sousa (2012) and Glick and Rose (2016) only provide cross-sectional estimates.

[48] As in Persson (2001) and Rose (2001), the probit regresses the currency union indicator on the product of the GDPs and the GDPs per capita, the log of distance, and dummy variables for sharing a common border, a common language, the same country, colonial relationships, and RTAs (available upon request). Due to the inclusion of GDP and GDP per capita, the probit is run on a smaller sample that includes 753,183 observations, of which 12,032 correspond to currency union pairs. Similar to Rose (2001), our results are unaffected by the value chosen for the maximum distance between the non-currency union observations and the propensity score. In Table C2, this distance is equal to 0.000001.

[49] The 12,032 currency union observations are matched with 57,781 non-currency union observations. The sample therefore includes a total of 69,813 observations, which represents one-ninth of our full sample size. The number of observations reported in Table C2 is smaller as the singletons are dropped due to the fixed effects.

Currency Union Types De Sousa (2012) provides information on currency union membership which we use to construct the indicator variable for common currencies (see the Data Appendix A). He identifies three types of currency unions: *multilateral* (i.e., between countries of similar size and wealth), *bilateral* (i.e., when a small or poor country adopts the currency of a larger and richer country), and currency unions where money is "interchangeable" between two countries at 1:1 parity. In Table C3 we broadly split currency unions into two groups, i.e., multilateral versus bilateral, and separately include a dummy variable for each group.[50] We then allow for heterogeneity in the trade impact of both multilateral and bilateral unions. As the interactions between the currency union dummy variables and predicted import shares are negative and significant, we conclude that within both types of unions the trade effects of common currencies are heterogeneous.

Currency Union Entry and Exit Our sample includes 342 and 459 (directional) switches into and out of currency unions. Among the 342 entries, 249 correspond to the euro.[51] To check whether the different types of switches matter for heterogeneity, we classify our currency union observations into three categories: *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). Some bilateral pairs are therefore classified both as entry and exit when they first entered and subsequently left a currency union.

Distinguishing between the three types of unions, we regress equation (9) and report the results in column (1) of Table C4 (for the continuous unions, the currency union dummy is omitted due to collinearity with the pair fixed effects). Interestingly, the interactions with the predicted shares are negative and significant for the continuous and exit unions only. In column (2), we further split the entry currency unions between euro and non-euro, and the interaction is negative and significant for the euro only.[52] We therefore conclude that with the exception of the non-euro entry currency unions, all the others are associated with heterogeneous trade effects, as predicted by the theoretical framework.

Import Shares per Good Our findings remain robust to using alternative proxies for the extensive margin $n_{i,t}$ in measuring the bilateral import shares per good. In columns (1), (2), and (3) of Table C5, respectively, the import shares per good are computed using the Hummels

[50] Interchangeable money observations are categorized according to the currency involved. The currencies used in multilateral unions include the British West Indies currency, the Central America and the Caribbean currency, the CFA and CFP francs, the East African currency, and the euro. The currencies circulating in bilateral unions are the Australian, Malaysian, and US dollars, the Indian, Mauritian, and Pakistani rupees, the Belgian and French francs, the South African rand, the Danish krone, the Portuguese escudo, the Saudi riyal, the Spanish peseta, and the British pound sterling.

[51] Belgium and Luxembourg are merged into a single entity, while Latvia and Lithuania only adopted the euro in 2014 and 2015, respectively. Our sample therefore includes 16 countries that switched to the euro, and they account for $16 \times 15 = 240$ directional switches. The nine other switches occurred between Saint Pierre et Miquelon and Eurozone countries.

[52] The regression includes a pair-specific trend for EU countries. The results remain similar without the trend.

and Klenow (2005) measure, the GDP of the exporting country, or assuming that the extensive margin is unity for all exporters.

Instead of using the importing country's GDP to compute the import shares per good, we experiment using total (column 4) or manufacturing (column 5) gross output from the OECD STAN database (available in domestic currency, and converted to US dollars using the bilateral exchange rates from the IMF's International Financial Statistics). As the data are only available for OECD nations, our sample is reduced to 19 importing countries, but our results continue to hold.

Specifications We consider two alternative specifications for the first-step regression (8) that generates the predicted import shares per good. First, in addition to bilateral distance and contiguity we include indicator variables for sharing a common language, a common colonizer post-1945, pairs in a colonial relationship post-1945, and for territories that were, or are, the same country.[53] Second, we replace the K_{ij} controls with a full set of (directional) country pair fixed effects. In both cases we derive the predicted import shares per good, which we then interact with the currency union dummy variable to estimate equation (9). The results based on the first step with the full set of gravity controls, and with the pair fixed effects, are reported in columns (1) and (2) of Table C6, respectively. Consistent with our baseline results, the interaction terms between the currency union dummy and the predicted shares are negative and highly significant. The trade impact of currency unions therefore falls with bilateral import intensity.

We also show that our results remain robust to including a lagged dependent variable (column 3), to explaining the log-difference of bilateral import shares between subsequent years (column 4), to controlling for a country pair-specific trend for EU countries (column 5), and to including more generally a country-pair specific trend for all countries in a currency union in our sample (column 6).

Samples Our main analysis uses the bilateral exports and GDP data from Head, Mayer, and Ries (2010), which we extended from 2007 to 2013 using the export flows and GDPs from the IMF's DOTS and the World Bank's WDI, respectively (see the Data Appendix A for more details). As a robustness check, we run our regressions using the original exports and GDP data from Head et al. (2010) over the 1949 to 2006 period, and the export data from the DOTS combined with the GDPs from the WDI between 1960 and 2013 (columns 1 and 2 of Table C7, respectively). These alternative samples leave our results qualitatively unchanged.

We also run regressions using a balanced sample over the period 1994 to 2013 (column 3), and at five-year intervals (column 4). We drop (in column 5) the countries (mostly island

[53] The estimated coefficients are significant and with the expected signs, import intensity being larger between closer and contiguous countries that share a common language, colonial ties, and are the same country. Our results are similar if we also control for the number of landlocked and island nations in each pair (but those are insignificant).

nations) omitted from the analysis of Glick and Rose (2016), the smaller nations with a nominal GDP less than 500 million US dollars in 2013 (column 6), and the poorer countries with an annual GDP per capita below 500 US dollars in 2013 (column 7). Finally, we restrict the sample to similarly sized country pairs, i.e., those with GDPs that differ less than threefold (column 8), and exclude the post-Soviet states (column 9).[54] Our results on heterogeneity hold up.

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. We then apply this framework to the effect of currency unions on international trade.

The prediction is that the impact of currency unions 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. The results lend strong support to our theoretical prediction. Our findings are robust to including zero import shares in the sample and allowing for non-random selection into currency unions. In line with the literature, we only find weak evidence that the euro has promoted bilateral trade on average among Eurozone members. However, we find that even within the Eurozone, the currency union effect is heterogeneous across and within country pairs.

By highlighting heterogeneity our results have implications for understanding some of the disparities in the currency union estimates reported in the literature. We agree with other researchers that the methodology (e.g., OLS versus PPML), the specification (e.g., types of fixed effects and right-hand side regressors), and the data used (e.g., Eurozone versus world-wide trade) matter for the magnitude and significance of currency union estimates. But our contribution is instead to show that relying on a single currency union estimate can be misleading if the objective is to assess, or to predict, the impact of sharing a common currency on bilateral trade flows. As we show, the trade effect of currency unions varies systematically, and it is larger for pairs associated with lower import intensity.

[54] The post-Soviet states are Armenia, Azerbaijan, Belarus, Estonia, Georgia, Kazakhstan, Kyrgyzstan, Latvia, Lithuania, Moldova, the Russian Federation, Tajikistan, Turkmenistan, Ukraine, and Uzbekistan.

References

Alesina, A., Barro, R.J., 2002. Currency Unions. *Quarterly Journal of Economics* 117(2), pp. 409–436.

Alesina, A., Barro, R.J., Tenreyro, S., 2002. Optimal Currency Areas. *NBER Macroeconomics Annual* 17, pp. 301–345.

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

Auboin, M., Ruta, M., 2013. The Relationship between Exchange Rates and International Trade: A Literature Review. *World Trade Review* 12(3), pp. 577–605.

Baier, S.L., Yotov, Y.V., Zylkin, T., 2016. On the Widely Differing Effects of Free Trade Agreements: Lessons from Twenty Years of Trade Integration. *LeBow College of Business, Drexel University School of Economics Working Paper Series* 2016–15.

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., 2006. Euros and Zeros: The Common Currency Effect on Trade in New Goods. *NBER Working Paper* 12673.

Baldwin, R., Di Nino, V., Fontagné, L., De Santis, R.A., 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., Taglioni, D., 2007. Trade Effects of the Euro: A Comparison of Estimators. *Journal of Economic Integration* 22(4), pp. 780–818.

Barro, R., Tenreyro, S., 2007. Economic Effects of Currency Unions. *Economic Inquiry* 45(1), pp. 1–23.

Berger, H., Nitsch, V., 2008. Zooming Out: The Trade Effect of the Euro in Historical Perspective. *Journal of International Money and Finance* 27(8), pp. 1244–1260.

Bergin, P.R., Lin, C.-Y., 2012. The Dynamic Effects of a Currency Union on Trade. *Journal of International Economics* 87(2), pp. 191–204.

Berthou, A., Fontagné, L., 2008. The Euro Effects on the Firm and Product-Level Trade Margins: Evidence from France. *CEPII Working Paper* 2008–21.

Broda, C., Romalis, J., 2010. Identifying the Relationship between Trade and Exchange Rate Volatility, in: Takatoshi, I., Rose, A.K. (Eds.), Commodity Prices and Markets, East Asia Seminar on Economics, Volume 20, University of Chicago Press, Chicago, pp. 79–110.

Bun, M.J.G., Klaassen, F.J.M., 2007. The Euro Effect on Trade is Not as Large as Commonly Thought. *Oxford Bulletin of Economics and Statistics* 69(4), pp. 473–496.

Campbell, D., 2013. Estimating the Impact of Currency Unions on Trade: Solving the Glick and Rose Puzzle. *The World Economy* 36(10), pp. 1278–1293.

Chaney, T., 2008. Distorted Gravity: The Intensive and Extensive Margins of International Trade. *American Economic Review* 98(4), pp. 1707–1721.

Correia, S., Guimarães, P., Zylkin, T., 2019. "ppmlhdfe": Fast Poisson Estimation with High-Dimensional Fixed Effects, mimeo.

De Sousa, J., 2012. The Currency Union Effect on Trade is Decreasing over Time. *Economics Letters* 117(3), pp. 917–920.

De Nardis, S., Vicarelli, C., 2003. Currency Unions and Trade: The Special Case of EMU. *Review of World Economics* 139(4), pp. 625–649.

Devereux, M.B., Lane, P.R., 2003. Understanding Bilateral Exchange Rate Volatility. *Journal of International Economics* 60(1), pp. 109–132.

Eaton, J., Kortum, S., 2002. Technology, Geography, and Trade. *Econometrica* 70(5), pp. 1741–1779.

Eicher, T.S., Henn, C., 2011. One Money, One Market: A Revised Benchmark. *Review of International Economics* 19(3), pp. 419–435.

Feenstra, R., 2003. A Homothetic Utility Function for Monopolistic Competition Models, without Constant Price Elasticity. *Economics Letters* 78(1), pp. 79–86.

Flam, H., Nordström, H., 2003. Trade Volume Effects of the Euro: Aggregate and Sector Estimates. *Institute for International Economic Studies Seminar Paper* 746, Stockholm University.

Flam, H., Nordström, H., 2007. Explaining Large Euro Effects on Trade: The Extensive Margin and Vertical Specialization. Institute for International Economic Studies, Stockholm University, mimeo.

Frankel, J., 2010. The Estimated Trade Effects of the Euro: Why are they below those from Historical Monetary Unions among Smaller Countries?, in: Alesina, A., Giavazzi, F. (Eds), Europe and the Euro. University of Chicago Press, pp. 169–212.

Frankel, J., Rose, A., 2002. An Estimate of the Effect of Common Currencies on Trade and Income. *Quarterly Journal of Economics* 117(2), pp. 437–466.

Frankel, J.A., Wei, S.-J., 1993. Trade Blocs and Currency Blocks. *NBER Working Paper* 4335.

Glick, R., 2016. Currency Unions and Regional Trade Agreements: EMU and EU Effects on Trade. *Federal Reserve Bank of San Francisco Working Paper* 2016–27.

Glick, R., Rose, A.K., 2002. Does a Currency Union Affect Trade? The Time-Series Evidence. *European Economic Review* 46(6), pp. 1125–1151.

Glick, R., Rose, A.K., 2016. Currency Unions and Trade: A Post-EMU Reassessment. *European Economic Review* 87, pp. 78–91.

Head, K., Mayer, T., 2014. Gravity Equations: Workhorse, Toolkit, and Cookbook, Chapter 3 in: Gopinath, G., Helpman, E., Rogoff, K. (Eds.), Handbook of International Economics, pp. 131–195.

Head, K., Mayer, T., Ries, J., 2010. The Erosion of Colonial Linkages after Independence. *Journal of International Economics* 81(1), pp. 1–14.

Hummels, D., Klenow, P.T., 2005. The Variety and Quality of a Nation's Exports. *American Economic Review* 95(3), pp. 704–723.

Irarrazabal, A., Moxnes, A., Opromolla, L.D., 2015. The Tip of the Iceberg: A Quantitative Framework for Estimating Trade Costs. *Review of Economics and Statistics* 97(4), pp. 777–792.

Klein, M.W., 2005. Dollarization and Trade. *Journal of International Money and Finance* 24(6), pp. 935–943.

Klein, M.W., Shambaugh, J.C., 2006. Fixed Exchange Rates and Trade. *Journal of International Economics* 70(2), pp. 359–383.

Larch, M., Wanner, J., Yotov, Y., Zylkin, T., 2017. The Currency Union Effect: A PPML Re-Assessment with High-Dimensional Fixed Effects. *LeBow College of Business Working Paper* 2017–07.

López-Córdova, J.E., Meissner, C.M., 2003. Exchange-Rate Regimes and International Trade: Evidence from the Classical Gold Standard Era. *American Economic Review* 93(1), pp. 344–353.

Machado, J.A.F., Santos Silva, J.M.C., Wei, K., 2016. Quantiles, Corners, and the Extensive Margin of Trade. *European Economic Review* 89, pp. 73–84.

Micco, A., Stein, E., Ordoñez, G., 2003. The Currency Union Effect on Trade: Early Evidence from EMU. *Economic Policy* 37(18), pp. 315–356.

Mika, A., Zymek, R., 2016. Friends without Benefits? New EMU Members and the "Euro Effect" on Trade. *Edinburgh School of Economics Discussion Paper* 269.

Mundell, R.A., 1961. A Theory of Optimum Currency Areas. *American Economic Review* 51(4), pp. 657–665.

Nitsch, V., 2002. Honey, I Shrunk the Currency Union Effect on Trade. *The World Economy* 25(4), pp. 457–474.

Novy, D., 2013. International Trade without CES: Estimating Translog Gravity. *Journal of International Economics* 89(2), pp. 271–282.

Parente, P.M.D.C., Santos Silva, J.M.C., 2016. Quantile Regression with Clustered Data. *Journal of Econometric Methods* 5(1), pp. 1–5.

Persson, T., 2001. Currency Unions and Trade: How Large is the Treatment Effect? *Economic Policy* 33(16), pp. 434–448.

Rose, A.K., 2000. One Money, One Market: Estimating the Effect of Common Currencies on Trade. *Economic Policy* 30(15), pp. 8–45.

Rose, A.K., 2001. Currency Unions and Trade: The Effect is Large. *Economic Policy* 16(33), pp. 449–461.

Rose, A.K., 2005. Which International Institutions Promote International Trade? *Review of International Economics* 13(4), pp. 682–698.

Rose, A.K., 2006. Currency Unions. *The New Palgrave Dictionary of Economics*. Second Edition. Steven N. Durlauf and Lawrence E. Blume (Eds). Palgrave Macmillan, 2008.

Rose, A.K., Stanley, T.D., 2005. A Meta-Analysis of the Effect of Common Currencies on International Trade. *Journal of Economic Surveys* 19(3), pp. 347–365.

Rose, A., van Wincoop, E., 2001. National Money as a Barrier to International Trade: The Real Case for Currency Union. *American Economic Review* 91(2), pp. 386–390.

Saia, A., 2017. Choosing the Open Sea: The Cost to the UK of Staying out of the Euro. *Journal of International Economics* 108, pp. 82–98.

Santos Silva, J.M.C., Tenreyro, S., 2006. The Log of Gravity. *Review of Economics and Statistics* 88(4), pp. 641–658.

Santos Silva, J.M.C., Tenreyro, S., 2010a. Currency Unions in Prospect and Retrospect. *Annual Review of Economics* 2, pp. 51–74.

Santos Silva, J.M.C., Tenreyro, S., 2010b. On the Existence of the Maximum Likelihood Estimates in Poisson Regression. *Economics Letters* 107(2), pp. 310–312.

Spearot, A.C., 2013. Variable Demand Elasticities and Tariff Liberalization. *Journal of International Economics* 89(1), pp. 26–41.

Subramanian, A., Wei, S.-J., 2007. The WTO Promotes Trade, Strongly but Unevenly. *Journal of International Economics* 72(1), pp. 151–175.

A Data Appendix

Bilateral Exports The IMF's Direction of Trade Statistics (DOTS) is the most widely used data set for studying the effect of currency unions on international trade (Alesina, Barro, and Tenreyro, 2002; Baldwin and Taglioni, 2007; Berger and Nitsch, 2008; Bun and Klaassen, 2007; De Nardis and Vicarelli, 2003; De Sousa, 2012; Glick, 2016; Glick and Rose, 2002, 2016; Micco, Stein, and Ordoñez, 2003; Mika and Zymek, 2016; Santos Silva and Tenreyro, 2010a). For more than 200 countries between 1948 and 2014, it reports their bilateral FOB merchandise exports (in current US dollars), of which 46 percent are recorded as being equal to zero. Head, Mayer, and Ries (2010) argue, however, that the true value of many of the zero export flows reported by the DOTS is likely to be positive. Relying on alternative data sources and on regression analysis, these authors identify a number of problematic zeros, and replace them by positive values or set them as missing entries. They also fix a number of typos that they argue are due to incorrect reporting between FOB versus CIF values. Our main analysis therefore relies on the data set cleaned by Head et al. (2010). But as their sample only spans the period from 1948 to 2006, we update their series up to 2014 using the growth rates of positive exports reported by the DOTS. As a robustness check, we also run regressions using the original export flows provided by the DOTS, and the data set compiled by Head et al. (2010) up to 2006.

GDP and Population Nominal GDP (in current US dollars) and population data are from Head et al. (2010) who complement, using national data sources and historical databases, the series from the World Bank's World Development Indicators (WDI) which start in 1960. As the series provided by Head et al. (2010) are available between 1949 and 2006 only, we update them up to 2013 using the growth rates of GDP and population from the WDI. GDP per capita is calculated as GDP divided by population.

Gravity Gravity controls are from the Centre d'Etudes Prospectives et d'Informations Internationales (CEPII). These include bilateral (population weighted) distances (in kilometers), and dummies for sharing a common land border (contiguity), a common (official) language, a common colonizer post-1945, pairs in a colonial relationship post-1945, and countries that were, or are, the same country. Dummy variables for membership with the OECD, IMF, and WTO are constructed using online sources (in each case, the dummy is equal to one if both countries in a pair are members in each year, and zero otherwise). De Sousa (2012) provides information on Regional Trade Agreements (RTAs) between 1958 and 2015. Using data from the CEPII we update his series from 1949 to 1957.

Currency Unions De Sousa (2012) provides information on currency union membership between 1948 and 2014 (based on Glick and Rose, 2002, and extended to include the euro). He identifies three types of currency unions: 1) *bilateral* currency unions, which "commonly occur when a small and/or poor country unilaterally adopts the money of a larger, richer 'anchor' country" (Rose, 2006), 2) *multilateral* currency unions "between countries of more or less equal size and wealth" (Rose, 2006), and 3) cases where "money was interchangeable between the two countries at a 1:1 par for an extended period of time, so that there was no need to convert prices when trading between a pair of countries."

De Sousa's (2012) data set includes 230 countries between 1948 and 2014 and reports 58,534 currency union observations. Between 1949 and 2013, which is the time period we focus on in our paper, this number drops to 54,648. As we describe in more detail below, our samples with and without the zero import share observations only include 19,514 and 13,085 currency union observations, respectively. There are several reasons for these differences. First, a number of currency union countries are omitted from the Head et al. (2010) data set on exports and GDPs. These include American Samoa, Belgium, Guam, Monaco, Luxembourg, and Montenegro (De Sousa, 2012, reports data for Belgium and Luxembourg both separately and as a single entity, while we aggregate them over the entire period). Second, for other countries the import shares per good cannot be calculated if either bilateral exports, the importer's GDP, or the extensive margin are not observed. These include currency union countries such as Montserrat, San Marino, and Wallis and Futuna which do not have any export and import data; the Falkland Islands, Gibraltar, Nauru, and Saint Helena which have no extensive margin and GDP data; Guadeloupe, French Guiana, Martinique, Réunion, and Saint Pierre et Miquelon which are omitted as importers as they have no GDP data; Andorra which is excluded as an importer because in the sample it only imports from Taiwan which does not have any extensive margin data; and finally, Equatorial Guinea which is omitted as an exporter because it lacks extensive margin data.

Note that other countries, which according to De Sousa (2012) never belonged to a currency union, are also excluded from our data set: Anguilla, the British Virgin Islands, Cocos Islands, Cook Islands, Christmas Island, Cayman Islands, Micronesia, Marshall Islands, Northern Mariana Islands, Norfolk Island, Niue, the Palestinian Territory, Pitcairn, Puerto Rico, Turks and Caicos Islands, Tokelau, and Western Sahara have no export and import data; North Korea, Taiwan, and Uzbekistan are excluded as exporters because of missing extensive margin data; Timor-Leste is excluded as an importer because in the sample it only imports from Taiwan which extensive margin is not available.

Descriptive Statistics As the pre-1997 trade flows for Belgium and Luxembourg are reported jointly, we aggregate the two countries into a single entity over the entire period (and count the two countries as one). Our main sample therefore includes 199 countries between 1949 and 2013. Bilateral import shares are given by the ratio between bilateral exports and the importing country's GDP, and we discard outliers by excluding the highest import shares that represent 0.05 percent of the sample size. Bilateral import shares *per good* are then obtained by dividing the import shares by the number of 4-digit HS- or SITC-level product categories exported by each country as a share of the total number of categories exported in each year, averaged over time (from United Nations Comtrade). In Section 3.1.1 we produce more details on how we measure the number of product categories.

As shown in Table A1, our full sample includes 1,203,583 observations, of which 782,469 import shares (and therefore import shares per good) are positive, and 421,114 are equal to zero (i.e., 35 percent of the sample). In the sample of positive import shares, the lowest import share is very close to zero percent (from Angola to Colombia), the largest is equal to 41.3 percent (from Singapore to the Maldives), and the mean and standard deviation are equal to 0.4 and 1.9 percent, respectively. As the import shares per good are given by the import shares over the extensive margin, they do not have any meaningful units and are therefore not described in Table A1. Overall, 924 country pairs (directional) share a common currency at least at some point (amounting to 13,085 observations, or about 1.7 percent of the positive import shares sample). There are 342 and 459 country pairs (directional) that switched into or out of currency unions, respectively.

Table A1: Descriptive Statistics

Full sample	
Total number of observations	1,203,583
Number of zero import shares	421,114
Number of positive import shares	782,469
Number of observations for currency unions	19,514
Positive import shares sample	
Import shares	
Minimum	0.001%
Maximum	41.264%
Mean	0.447%
Standard deviation	1.869%
Number of observations for currency unions	13,085
Number of pairs in a currency union (directional)	924
Number of switches into currency unions (directional)	342
Number of switches out of currency unions (directional)	459

Source: Authors' calculations.

Table A2 provides descriptive statistics for currency unions and non-unions in the sample of positive import shares. For most variables, the sample means are similar for both groups of countries, with some exceptions: countries in a currency union have higher import shares, are closer, are more likely to be in a colonial relationship and more likely to belong to both the OECD and WTO.

Table A2: Currency Unions and Non-Unions

	Non-Unions	Currency Unions
Import share (%)	0.427 (1.777)	1.583 (4.672)
RTA	0.071 (0.256)	0.399 (0.490)
IMF	0.801 (0.398)	0.802 (0.398)
OECD	0.039 (0.193)	0.147 (0.355)
WTO	0.519 (0.500)	0.685 (0.464)
ln Distance	8.617 (0.820)	7.411 (0.929)
Contiguity	0.027 (0.163)	0.184 (0.387)
Shared language	0.168 (0.374)	0.719 (0.449)
Common colonizer	0.083 (0.275)	0.563 (0.496)
Colonial relationship	0.013 (0.115)	0.069 (0.253)
Same country	0.010 (0.101)	0.210 (0.407)
Observations	769,384	13,085

Notes: The table reports the mean and standard deviation (in parentheses) of each variable.

B Monte Carlo Analysis

B.1 Analysis of the Two-Step Procedure

When running log-linear gravity regressions in Section 3, we adopted a two-step procedure to estimate heterogeneous currency union effects. In the first step, we predicted the import shares per good. In the second step, we interacted the currency union dummy with the predicted shares. In this section we carry out Monte Carlo simulations to verify the validity of this two-step procedure.

As our trade cost function, we assume:

$$\ln(t_{ij,t}) = \kappa CU_{ij,t} + \zeta W_{ij,t}, \tag{B1}$$

where $W_{ij,t}$ contains all bilateral trade cost variables used in our analysis, i.e., time-invariant geography-related variables (logarithmic bilateral distance and a contiguity dummy) as well as time-varying policy variables (dummies for RTAs and membership of the IMF, OECD, and WTO). We choose values for the trade cost parameters that are derived from our baseline regression in column (2) of Table 1.[55] We then compute trade costs on the basis of equation (B1) using the actual observations for our trade cost variables.

We assume that the true data generating process is given by the translog gravity model in Section 2. We choose the translog parameter value as $\theta = 0.073$.[56] Based on equations (1)–(4), we first compute the import shares in a deterministic way (i.e., without an error term). We use a balanced sample of observed data for the GDP variables ($y_{i,t}, y_{j,t}$) and the trade cost variables underlying equation (B1) for 120 origin and destination countries over the period from 1990 to 2013.[57]

Then we include an additive error term in the translog gravity equation (1). We choose its standard deviation to match the R-squared of around 60 percent in our translog regressions of Table 5.[58] We take the natural logarithm of the simulated import shares, thus dropping non-positive values. We run first-step and second-step regressions as in Section 3 where the import shares predicted in the first step are interacted with the currency union dummy in the second step. Standard errors are clustered by country pairs. For simplicity we assume $n_{i,t} = 1$

[55] Assuming an elasticity of substitution of $\sigma = 5$ for the standard log-linear gravity framework as in Anderson and van Wincoop (2003), the κ parameter for the currency union dummy in (B1) follows as the estimated coefficient of 0.326 in column (2) of Table 1 divided by $(1-\sigma)$, i.e., $\kappa = 0.326/(1-5) = -0.082$. The parameters for RTA, IMF, OECD and WTO follow analogously as -0.104, -0.041, -0.092, and -0.037, respectively. For distance and contiguity, we run a regression as in (8) based on the observed import shares, with estimated coefficients of -1.411 and 0.805 (both significant at the one percent level). Their parameters in (B1) thus follow as 0.353 and -0.201.

[56] In the translog regression in column (1) of Table 5, we obtain a currency union coefficient of 0.006. Assuming the same currency union coefficient as above, it therefore follows $\theta = -0.006/\kappa = 0.073$.

[57] To reduce computing time we use a subset of data starting in 1990.

[58] A value of 0.028 for the standard deviation leads to an R-squared of 62 percent.

for all i.[59] We run 100 iterations of this procedure, drawing a new set of error terms for every iteration.

Table B1: Monte Carlo Simulation

	(1)	(2)	(3)
		First step	Second step
CU	–	–	-1.103^{a} (0.168)
CU×predicted share	–	–	-0.363^{a} (0.049)
RTA	–	–	0.145^{a} (0.017)
IMF	–	–	0.085 (0.162)
OECD	–	–	0.141^{a} (0.042)
WTO	–	–	0.049^{c} (0.030)
ln Distance	–	-0.546^{a} (0.005)	–
Contiguity	–	0.226^{a} (0.014)	–
CU estimates	**True**		**Estimated**
Mean	0.278	–	0.290^{a} (0.049)
10^{th} percentile	1.262	–	0.839^{a} (0.107)
90^{th} percentile	0.080	–	-0.162^{a} (0.056)
R-squared	–	0.307	0.373
Observations	–	187,468	187,468
Corresponding table, (column)	–	–	Table 1, (4)

Notes: Exporter-year and importer-year fixed effects are included. Directional country pair fixed effects are further included in (3). The regressions are estimated between 1990 and 2013 based on OLS estimation of a log-linear gravity model with the log import share as the dependent variable. But translog gravity is the true underlying data generating process (with the true currency union elasticities reported in the first column). The reported coefficients are averages over the 100 iterations. Robust standard errors clustered at the country pair level. [a] and [c] indicate significance at the one and ten percent levels, respectively. "predicted share" is the predicted log import share.

We report the results in Table B1, averaged over all iterations. Analogous to specification (8), the first-step regression in column (2) simply includes distance and a contiguity dummy, and it yields the expected signs. The second-step regression in column (3) includes the currency union dummy and an interaction term with the log predicted import shares, as well as the additional time-varying policy variables. Consistent with column (4) of Table 1, we obtain negative coefficients on both the currency union dummy and the interaction term. The lower panel of column (3) reports the implied currency union estimates, evaluated at the mean, 10^{th}, and the 90^{th} percentiles of import shares. We find a mean estimate of 0.290, implying that evaluated at the average import share, two countries trade 33.6 percent more bilaterally if they

[59]We also ran specifications with the extensive margin measure as observed in the data based on Comtrade (see Section 3.1.1). The overall results are very similar.

are in a currency union. Consistent with the theoretical framework, we find a larger estimate of 0.839 at the 10^{th} percentile (i.e., for relatively small import shares), implying 131 percent more bilateral trade ceteris paribus. At the 90^{th} percentile (i.e., for relatively large import shares) we find an estimate of -0.162, implying reduced bilateral trade by 15 percent. These estimates can be compared to the true values underlying the simulation indicated in column (1).

Overall, the simulation in Table B1 confirms the validity of our two-step procedure in the sense that qualitatively, it yields the same results as in the underlying true model. Quantitatively, our results are similar for the mean, but they somewhat undershoot the true effect at the 10^{th} and 90^{th} percentiles. Figure 1 visually compares the true values against the estimates across the deciles of (predicted) import shares. 95 percent confidence intervals are indicated in dashed lines. The true values lie within the confidence intervals – except for very small and very large percentiles. The reason for the relatively large deviation between the true and the estimated values at the lowest percentile is the functional form of the translog specification. As equation (5) shows, the translog elasticity is given by the translog preference parameter divided by the import share. This generates a hyperbolic shape such that elasticity values decline rapidly with growing import shares.[60] Our simple interaction term between the currency union dummy and the log predicted import shares struggles to capture the very large effects at the smallest import shares, but it matches the remaining percentiles rather well.

As an additional check, we also investigate the consequences of ignoring the first step altogether by erroneously interacting the currency union dummy with *actual* log import shares (as opposed to predicted log import shares). Since in that case the interacted regressor is by construction positively correlated with the dependent variable, this leads to an upward endogeneity bias on the interaction coefficient. In fact, it even turns positive with high significance.[61] The resulting estimates at the mean, 10^{th} and the 90^{th} percentiles follow as -0.436, -1.860, and 0.735 (all significant at the one percent level). Thus, they exhibit the opposite pattern of the true values in Table B1 in that they *rise* with the import share, which is incorrect. This check therefore underlines the importance of predicting shares in the first step.

B.2 Placebo Checks

We also carry out placebo checks that are based on the assumption that the standard log-linear gravity model represents the true underlying data generating process. We construct the import shares for the standard gravity model using the relationship:

$$\frac{x_{ij,t}}{y_{j,t}} = \frac{y_{i,t}}{y_t^W} \left(\frac{t_{ij,t}}{P_{i,t} P_{j,t}} \right)^{1-\sigma}, \tag{B2}$$

[60] Also see the discussion in Section 3.2.

[61] This pattern is also confirmed by our regressions based on observed data. If in equation (9) we interact the currency union dummy with actual import shares per good, the interaction term is positive and significant at the one percent level.

Table B2: Monte Carlo Simulation (Placebo Checks)

	(1)	(2)	(3)
		First step	Second step
CU	–	–	0.329^a (0.114)
CU×predicted share	–	–	0.001 (0.019)
RTA	–	–	0.418^a (0.013)
IMF	–	–	0.180 (0.150)
OECD	–	–	0.368^a (0.035)
WTO	–	–	0.143^a (0.024)
ln Distance	–	-1.533^a (0.005)	–
Contiguity	–	0.964^a (0.023)	–
CU estimates	**True**		**Estimated**
Mean	0.326	–	0.323^a (0.058)
10^{th} percentile	0.326	–	0.320^a (0.111)
90^{th} percentile	0.326	–	0.326^a (0.048)
R-squared	–	0.811	0.822
Observations	–	345,600	345,600
Corresponding table, (column)	–	–	Table 1, (4)

Notes: Exporter-year and importer-year fixed effects are included. Directional country pair fixed effects are further included in (3). The regressions are estimated between 1990 and 2013 based on OLS estimation of a standard log-linear gravity model with the log import share as the dependent variable. Standard gravity is the true underlying data generating process (with the true currency union elasticities reported in the first column). The reported coefficients are averages over the 100 iterations. Robust standard errors clustered at the country pair level. [a] indicates significance at the one percent level. "predicted share" is the predicted log import share.

which is derived by Anderson and van Wincoop (2003). $P_{i,t}$ and $P_{j,t}$ denote the price indices of the origin and destination countries, or multilateral resistance terms, given by:

$$P_{i,t}^{1-\sigma} = \sum_{s=1}^{S} P_{s,t}^{\sigma-1} \frac{y_{s,t}}{y_t^W} t_{si,t}^{1-\sigma}, \tag{B3}$$

where S is the number of countries in the world. We assume $\sigma = 5$. We use equations (B2)–(B3) as well as trade cost function (B1) to construct the deterministic import shares, based on the same sample of GDP and trade cost variables for 120 countries as above. We solve for the price indices numerically through iteration. We choose the standard deviation of a lognormal multiplicative error term in order to match the R-squared of close to 81 percent in our baseline gravity regressions (see Table 1). As above, we run first-step and second-step regressions, iterating the procedure 100 times with fresh error terms.

We report the results of the placebo checks in Table B2. The first-step regression in column

(2) includes coefficients on distance and contiguity with the expected signs and magnitudes. Most importantly, the second-step regression in column (3) does not exhibit a significant interaction term for the currency union dummy and the log predicted import shares. Therefore, the estimated currency union effects reported in the lower panel do not vary between the mean, the 10^{th}, and the 90^{th} percentiles. Neither are they significantly different from the true effects in column (1).

Overall, the placebo results confirm that if standard gravity were the true underlying model, our two-step framework would not give rise to currency union effects that systematically differ across percentiles. We therefore conclude that it is important to use a gravity model that accommodates variable elasticities to capture the heterogeneous patterns we find in our main regression results.

B.3 General Equilibrium Effects

Our results on the heterogeneity of currency union effects reported in Section 3 refer to the *direct* effect of trade costs on trade (also see our discussion of the trade cost elasticity in Section 2). However, a change in trade costs also has an *indirect* effect on trade through multilateral resistance effects in general equilibrium, a point famously made by Anderson and van Wincoop (2003). The aim of this appendix is to trace out these general equilibrium effects in response to a change in trade costs. We show that they cannot explain the heterogeneity patterns in Section 3.

We base our analysis on the standard gravity model as in equation (B2). Similar to Novy (2013) we take logarithms and first differences to arrive at:

$$\Delta \ln \left(\frac{x_{ij,t}}{y_{j,t}} \right) = (1-\sigma)\,\Delta \ln \left(t_{ij,t} \right) + (\sigma - 1)\,\Delta \ln \left(P_{i,t} P_{j,t} \right) + \Delta \ln \left(\frac{y_{i,t}}{y_t^W} \right). \tag{B4}$$

The left-hand side represents the percentage change in the import share. The first term on the right-hand side represents the direct effect of the change in trade costs. The second and third terms indicate the indirect general equilibrium effects, consisting of the change in multilateral resistance and the change in the exporting country's income share.

With the help of decomposition (B4) we analyze a counterfactual change in trade costs. As in Section B.1 we draw on a sample of 120 origin and destination countries, using the observed data for the GDP variables ($y_{i,t}$, $y_{j,t}$) and the trade cost variables underlying equation (B1) with the same parameter values as in that section. In particular, the value of the currency union dummy coefficient κ is chosen such that it matches the 0.326 coefficient in column (2) of Table 1. Based on those data we numerically compute an initial equilibrium under the assumption of no currency unions (i.e., we set $CU_{ij,t}$ to zero for all pairs). As our counterfactual exercise, we then compute a new equilibrium under the assumption that a currency union is formed for a particular pair (i.e., we set $CU_{ij,t}$ to one for that particular pair). We then compute the terms in decomposition (B4), assuming that the exporting country's income share is constant.[62] We compute such a counterfactual equilibrium for each of the currency union pairs in our sample. We use data for a single-cross section (for the year 2000), in which the data indicate 244 currency union pairs. Thus, we compute 244 counterfactual equilibria.

Table B3: General Equilibrium Effects

Import share interval	Total effect $\Delta \ln (x_{ij}/y_j)$ (1)		Direct effect $(1-\sigma)\Delta \ln (t_{ij})$ (2)	Indirect GE effect $+(\sigma-1)\Delta \ln (P_i P_j)$ (3)
First interval	0.317	=	0.326	−0.009
Second interval	0.302	=	0.326	−0.024
Third interval	0.328	=	0.326	+0.002
Fourth interval	0.316	=	0.326	−0.010

Notes: This table is based on the decomposition in equation (B4). It reports logarithmic differences between the initial equilibrium and counterfactual equilibria (computed numerically). The initial equilibrium assumes no currency unions. In the counterfactual equilibria the currency unions are activated separately for each pair. The results are reported as averages over currency union pairs by quartiles of import shares, where intervals are formed in ascending order of import shares based on the initial equilibrium. Data are for 120 origin and destination countries in the year 2000 consisting of 244 currency union pairs (see Section B.1 for details including underlying parameter values). Column (1) reports the change in import shares, (2) reports the direct effect of entering a currency union, and (3) reports the indirect general equilibrium (GE) effect operating through multilateral resistance. Income shares are held constant.

We present the results in Table B3. Since we are interested in variation across import shares, we report the results as averages by import share intervals. Specifically, we choose four import share intervals in ascending order based on the initial equilibrium. For example, the first row of Table B3 reports the average changes for currency union pairs in the quartile of the smallest import shares in the initial equilibrium. By construction the direct effect in column (2) reflects the 0.326 currency union dummy coefficient from column (2) of Table 1. That is, entering into a currency union is associated with an increase in bilateral trade of 38 percent (equal to $\exp(0.326) - 1$). The indirect effect operating through changes in multilateral resistance in column (3) is quantitatively small. The intuition is that currency unions are relatively rare at the bilateral level (see the Data Appendix A), and they are only one out of several trade cost components. Therefore, the total effect in column (1) is similar to the direct effect.

Furthermore, we note that the indirect effect does not vary systematically across import share intervals. Intuitively, in response to a change in bilateral trade costs multilateral resistance typically shifts more strongly for small countries as they tend to be more open. But currency union pairs in the data are associated with a mix of both small and large countries across all intervals. Therefore, multilateral resistance effects do not vary systematically across

[62] That is, $\Delta \ln \left(y_{i,t}/y_t^W\right) = 0$. The effect operating through a changing income share is typically negligibly small.

intervals.[63]

In summary, we conclude that indirect trade cost effects tend to be quantitatively weak. They do not vary systematically across import share intervals. We also refer to Novy (2013, Section 3.5) who shows formally in Monte Carlo simulations that general equilibrium effects would in any case be absorbed by exporter and importer fixed effects in gravity regressions. Thus, heterogeneity patterns of the type we find in Section 3 are not related to general equilibrium effects.

[63] In a second counterfactual experiment (not reported here), we assume that the currency union dummy is set to one for *all* pairs at the same time. Quantitatively, the indirect general equilibrium effects are larger than in column (3) of Table B3 since this counterfactual experiment involves multiple trade cost changes in one go. But as in Table B3, there is no systematic variation across import share intervals.

C Robustness Appendix

Table C1: Robustness: Zero Trade Observations

	(1)	(2)	(3)	(4)	(5)	
CU	0.202^{a} (0.050)	-0.340^{a} (0.130)	0.252^{a} (0.055)	-0.378^{b} (0.149)	–	
CU×predicted share	–	-0.175^{a} (0.035)	–	-0.196^{a} (0.036)	–	
CU non EURO	–	–	–	–	-0.392^{c} (0.201)	
CU non EURO×predicted share	–	–	–	–	-0.216^{a} (0.046)	
EURO	–	–	–	–	-0.562^{a} (0.151)	
EURO×predicted share	–	–	–	–	-0.118^{a} (0.036)	
Trend EU countries	–	–	–	–	0.026^{a} (0.003)	
CU estimates					Non EURO	EURO
Mean	–	0.826^{a} (0.122)	–	0.979^{a} (0.129)	1.106^{a} (0.170)	0.259^{b} (0.130)
10^{th} percentile	–	1.251^{a} (0.202)	–	1.504^{a} (0.219)	1.685^{a} (0.282)	0.577^{a} (0.219)
90^{th} percentile	–	0.412^{a} (0.058)	–	0.477^{a} (0.065)	0.552^{a} (0.096)	-0.044 (0.069)
Zeros included	No	No	Yes	Yes	Yes	
Observations	780,818	780,818	1,131,641	1,131,641	1,131,641	

Notes: PPML estimations where exporter-year, importer-year, and (directional) country pair fixed effects are included. Robust standard errors adjusted for clustering at the (non-directional) country pair level are reported in parentheses in (1) and (3). Standard errors are boostrapped in (2), (4), and (5). ^{a}, ^{b}, and ^{c} indicate significance at the one, five, and ten percent levels, respectively. The dependent variable is the import share per good in levels. "predicted share" is the predicted import share per good. Dummy variables for RTAs, IMF, OECD, and WTO memberships are included but not reported.

Table C2: Robustness: Non-Random Selection

	(1)	(2)
CU	0.352^{a} (0.081)	-0.161 (0.258)
CU×predicted share	–	-0.086^{b} (0.043)
CU estimates		
Mean	–	0.509^{a} (0.124)
10^{th} percentile	–	0.815^{a} (0.257)
90^{th} percentile	–	0.214^{b} (0.102)
R-squared	0.904	0.904
Observations	64,779	64,779

Notes: Exporter-year, importer-year, and (directional) country pair fixed effects are included. Robust standard errors adjusted for clustering at the (non-directional) country pair level are reported in parentheses in (1). Standard errors are bootstrapped in (2). [a] and [b] indicate significance at the one and five percent levels, respectively. The dependent variable is the log import share per good. The sample includes the currency union observations and the non-currency union observations that deviate by no more than a small distance (equal to 0.000001) from the propensity score estimated from a probit regression of the currency union dummy on a number of gravity regressors. "predicted share" is the predicted log import share per good. Dummy variables for RTAs, IMF, OECD, and WTO memberships are included but not reported.

Table C3: Robustness: Currency Union Types

	(1)	
CU multilateral	0.040 (0.117)	
CU multilateral×predicted share	-0.052^{b} (0.021)	
CU bilateral	-0.317 (0.241)	
CU bilateral×predicted share	-0.091^{a} (0.034)	
CU estimates	Multilateral	Bilateral
Mean	0.469^{a} (0.095)	0.435^{a} (0.094)
10^{th} percentile	0.641^{a} (0.156)	0.735^{a} (0.183)
90^{th} percentile	0.299^{a} (0.060)	0.137 (0.097)
R-squared	0.808	
Observations	780,818	

Notes: Exporter-year, importer-year, and (directional) country pair fixed effects are included. Bootstrapped standard errors adjusted for clustering at the (non-directional) country pair level are reported in parentheses. [a] and [b] indicate significance at the one and five percent levels, respectively. The dependent variable is the log import share per good. "predicted share" is the predicted log import share per good. Dummy variables for RTAs, IMF, OECD, and WTO memberships are included but not reported.

Table C4: Robustness: Currency Union Entry and Exit

	(1)		(2)		
CU entry	0.442^a		–		
	(0.103)				
CU entry×predicted share	0.023		–		
	(0.020)				
CU non EURO entry	–		−0.435		
			(0.590)		
CU non EURO entry×predicted share	–		−0.054		
			(0.074)		
EURO entry	–		−0.198^b		
			(0.099)		
EURO entry×predicted share	–		−0.044^b		
			(0.020)		
CU exit	−0.405		−0.311		
	(0.249)		(0.265)		
CU exit×predicted share	−0.098^a		−0.087^b		
	(0.033)		(0.038)		
CU continuous×predicted share	−0.317^a		−0.307^a		
	(0.075)		(0.072)		
Trend EU countries	–		0.028^a		
			(0.003)		
CU estimates	Entry	Exit	Non EURO entry	EURO entry	Exit
Mean	0.254^a	0.399^a	0.009	0.163	0.405^a
	(0.096)	(0.088)	(0.216)	(0.110)	(0.095)
10^{th} percentile	0.178	0.721^a	0.187	0.308^c	0.691^a
	(0.155)	(0.168)	(0.361)	(0.168)	(0.200)
90^{th} percentile	0.328^a	0.081	−0.167	0.021	0.122
	(0.057)	(0.106)	(0.280)	(0.069)	(0.098)
R-squared	0.808		0.808		
Observations	780,818		780,818		

Notes: Exporter-year, importer-year, and (directional) country pair fixed effects are included. Bootstrapped standard errors adjusted for clustering at the (non-directional) country pair level are reported in parentheses. ^{a}, ^{b}, and ^{c} indicate significance at the one, five, and ten percent levels, respectively. The dependent variable is the log import share per good. "predicted share" is the predicted log import share per good. Dummy variables for RTAs, IMF, OECD, and WTO memberships are included but not reported.

Table C5: Robustness: Import Shares

	(1)	(2)	(3)	(4)	(5)
CU	0.053	0.080	0.029	-0.671^{c}	-0.344
	(0.109)	(0.086)	(0.102)	(0.346)	(0.266)
CU×predicted share	-0.069^{a}	-0.097^{a}	-0.045^{a}	-0.149^{b}	-0.135^{b}
	(0.021)	(0.022)	(0.015)	(0.064)	(0.057)
CU estimates					
Mean	0.524^{a}	0.574^{a}	0.428^{a}	0.612^{a}	0.625^{a}
	(0.076)	(0.069)	(0.065)	(0.235)	(0.173)
10^{th} percentile	0.699^{a}	0.817^{a}	0.601^{a}	1.253^{b}	1.211^{a}
	(0.119)	(0.112)	(0.113)	(0.504)	(0.409)
90^{th} percentile	0.345^{a}	0.302^{a}	0.263^{a}	-0.119	-0.034
	(0.056)	(0.056)	(0.049)	(0.128)	(0.147)
Exporter extensive margin	HK (2005)	GDP	Unity	Comtrade	Comtrade
Importer output	GDP	GDP	GDP	Total output	Manuf. output
R-squared	0.771	0.769	0.840	0.921	0.917
Observations	623,141	760,883	795,189	73,480	89,040

Notes: Exporter-year, importer-year, and (directional) country pair fixed effects are included. Bootstrapped standard errors adjusted for clustering at the (non-directional) country pair level are reported in parentheses. ^{a}, ^{b}, and ^{c} indicate significance at the one, five, and ten percent levels, respectively. The dependent variable is the log import share per good. "predicted share" is the predicted log import share per good. In (1), HK (2005) stands for Hummels and Klenow (2005). Dummy variables for RTAs, IMF, OECD, and WTO memberships are included but not reported.

Table C6: Robustness: Specifications

	(1)	(2)	(3)	(4)	(5)	(6)
Lagged dep. var.	–	–	0.555^a (0.003)	–	–	–
CU	-0.010 (0.119)	-0.231^c (0.134)	-0.013 (0.056)	-0.058^b (0.026)	-0.453^a (0.113)	-0.044 (0.109)
CU×predicted share	-0.059^a (0.019)	-0.093^a (0.022)	-0.030^a (0.009)	-0.010^b (0.005)	-0.103^a (0.017)	-0.052^a (0.017)
Trend EU countries	–	–	–	–	0.028^a (0.003)	–
Trend CU pairs	–	–	–	–	–	0.005 (0.004)
CU estimates						
Mean	0.473^a (0.076)	0.532^a (0.075)	0.224^a (0.033)	0.022 (0.017)	0.391^a (0.070)	0.384^a (0.091)
10^{th} percentile	0.668^a (0.128)	0.892^a (0.152)	0.319^a (0.056)	0.054^c (0.030)	0.729^a (0.112)	0.555^a (0.132)
90^{th} percentile	0.281^a (0.058)	0.202^a (0.052)	0.131^a (0.025)	-0.009 (0.010)	0.057 (0.058)	0.214^a (0.071)
Dependent variable	Log	Log	Log	Log-diff.	Log	Log
First-step pair controls	Gravity	Pair FE	Dist/Contig	Dist/Contig	Dist/Contig	Dist/Contig
R-squared	0.808	0.808	0.879	0.083	0.808	0.808
Observations	780,818	780,818	702,693	702,693	780,818	780,818

Notes: Exporter-year, importer-year, and (directional) country pair fixed effects are included. Bootstrapped standard errors adjusted for clustering at the (non-directional) country pair level are reported in parentheses. ^{a}, ^{b}, and ^{c} indicate significance at the one, five, and ten percent levels, respectively. The dependent variable is the log import share per good in all columns but the log-difference of the import share per good in (4). "predicted share" is the predicted log import share per good. Dummy variables for RTAs, IMF, OECD, and WTO memberships are included but not reported.

Table C7: Robustness: Samples

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
CU	−0.165	−0.075	0.102	−0.043	0.014	0.129	0.062	−0.062	0.022
	(0.144)	(0.117)	(0.119)	(0.125)	(0.129)	(0.118)	(0.117)	(0.287)	(0.116)
CU×predicted share	-0.069^a	-0.065^a	-0.038^c	-0.072^a	-0.057^a	-0.042^b	-0.047^a	−0.073	-0.054^a
	(0.021)	(0.018)	(0.023)	(0.022)	(0.021)	(0.019)	(0.018)	(0.046)	(0.018)
CU estimates									
Mean	0.402^a	0.480^a	0.383^a	0.548^a	0.482^a	0.462^a	0.433^a	0.536^a	0.465^a
	(0.071)	(0.073)	(0.080)	(0.094)	(0.076)	(0.067)	(0.068)	(0.134)	(0.076)
10^{th} percentile	0.624^a	0.707^a	0.500^a	0.785^a	0.670^a	0.591^a	0.577^a	0.746^a	0.644^a
	(0.123)	(0.122)	(0.140)	(0.157)	(0.131)	(0.113)	(0.110)	(0.247)	(0.122)
90^{th} percentile	0.182^a	0.254^a	0.271^a	0.313^a	0.296^a	0.332^a	0.288^a	0.334^a	0.289^a
	(0.063)	(0.057)	(0.053)	(0.060)	(0.058)	(0.053)	(0.058)	(0.094)	(0.061)
Sample	Head et al.	IMF	Balanced	5-year	Excl.	Excl. small	Excl. poor	Similar size	Excl. Soviet
	(2010)	DOTS	1994–2013	intervals	islands	countries	countries	countries	countries
R-squared	0.815	0.832	0.887	0.824	0.808	0.806	0.811	0.804	0.808
Observations	644,439	607,218	237,340	158,281	695,467	651,305	609,695	176,799	717,819

Notes: Exporter-year, importer-year, and (directional) country pair fixed effects are included. Bootstrapped standard errors adjusted for clustering at the (non-directional) country pair level are reported in parentheses. ^{a}, ^{b}, and ^{c} indicate significance at the one, five, and ten percent levels, respectively. The dependent variable is the log import share per good. “predicted share” is the predicted log import share per good. Dummy variables for RTAs, IMF, OECD, and WTO memberships are included but not reported.