

Friends Without Benefits?

New EMU Members and the “Euro Effect” on Trade

Alina Mika*
European Central Bank

Robert Zymek†
University of Edinburgh
and CESifo

October 2017

Abstract

We re-visit the evidence about the trade benefits of European Monetary Union (EMU), focusing on the experience of countries which adopted the common currency since 2002. Based on “state of the art” gravity estimations for the period 1992-2013, we reach three main conclusions. First, estimates from an appropriately specified and estimated gravity equation provide no evidence of a euro effect on trade flows among early euro adopters up to the year 2002. Second, this finding is robust to extending the sample period to incorporate data up to 2013, covering five additional euro accessions. Third, while there is no robust evidence of a euro effect, there is evidence that intra-EU trade flows have expanded faster than the global average during the 2002-2013 period. Using the functional form of a theory-consistent gravity equation, we perform pseudo out-of-sample forecasts of trade flows for recent euro joiners. In line with our estimation results, we show that pseudo forecasts of the change in trade flows after euro accession, assuming no euro effect, outperform forecasts based on the expectation of a significantly positive effect. This suggests that euro accession countries should not expect a significant boost to their trade from joining EMU.

JEL Classification codes: F14, F15, F17, F33

Keywords: euro, trade, gravity, poisson

*European Central Bank, Sonnemanstraße 20, 60314 Frankfurt am Main, Germany; alina.mika@ecb.int.

†School of Economics, University of Edinburgh, 31 Buccleuch Place, Edinburgh, EH8 9JT, United Kingdom; robert.zymek@ed.ac.uk.

We are grateful to Mariko Klasing, Petros Milionis, Andrew Rose, Silvana Tenreyro, and an anonymous referee for helpful comments. The views expressed here are those of the authors and should not be reported as representing the views of the European Central Bank (ECB).

1 Introduction

Since 2007, the eurozone has gained seven new members: Slovenia, Cyprus, Malta, Slovakia, Estonia, Latvia and Lithuania. Seven more countries are expected to join eventually.¹ The question regarding the economic benefits and costs to be expected from euro membership remains alive and well in these countries as they weigh the timing of their accession. In this paper, we re-visit the evidence about one of the benefits of European Monetary Union (EMU): increased trade integration among EMU members. Earlier empirical studies suggest that euro adoption resulted in a significant positive boost to trade flows among eurozone economies. However, these studies were based on early-days data from the original euro club, comprised of Western and Southern European countries, which had adopted the euro by 2002. Our paper focuses on the experience of subsequent joiners. In particular, we ask whether initial estimates of the “euro effect” on trade flows were a good guide to the euro’s impact on trade for later additions to the eurozone. Our findings imply that the answer is no – there is no robust evidence of a euro effect on trade flows, for recent as well as original adopters of the common currency.

[Insert Figure 1 about here]

Figure 1 provides a graphical preview of our results. For all countries which joined the euro between 2002 and 2013, it plots the value of their trade with the original euro members as a share of the value of their trade with all EU countries.² The year prior to their accession is used as the base year, and their accession year is marked with a vertical line. Estimates from earlier studies would lead us to expect an increase in trade flows of 5-15% with euro countries upon accession, holding everything else constant (see Baldwin, 2006). In the figure, there is no systematic evidence of such a rise in countries’ trade flows with the core eurozone – relative to the wider EU – in the wake of their euro adoption. In the remainder of the paper, we will confirm this casual observation using formal econometric methods.

Like earlier studies on the euro effect, we assess the impact of euro adoption on trade flows by estimating a gravity equation. Unlike these studies, we are able to make use of advances in computer processing power to employ a near-comprehensive data set of bilateral trade flows in the period 1992-2013, and to mitigate the risk of omitted-variable bias by using a full array of country-pair and country-time fixed effects. We also estimate our preferred gravity equation in levels using the Poisson pseudo-maximum likelihood (PPML) procedure recommended by Santos Silva and Tenreyro (2006).³ Our gravity estimation thus represents the current “state of the art” in the literature. After providing an updated estimate of the impact of euro adoption on trade, we

¹Apart from Denmark and the United Kingdom, all current EU member states are legally obliged to join the eurozone. At present, this applies to Bulgaria, Croatia, the Czech Republic, Hungary, Poland, Romania and Sweden.

²The data on the value of bilateral trade flows used to construct this figure is taken from the latest edition of the IMF’s *Direction of Trade Statistics*.

³Santos Silva and Tenreyro (2006) show that the traditional estimation of a log-linearised gravity equation using OLS will result in biased coefficients in the presence of a heteroskedastic error term, and Santos Silva and Tenreyro (2010) document that this issue may have led to biases in earlier estimations of the trade-effect of currency unions.

perform pseudo out-of-sample forecasts of trade flows for recent euro joiners using the functional form of a theory-consistent gravity equation. In this way, we can compare the accuracy with which different estimates of the euro effect would have predicted the post-accession evolution of trade for these countries.

Our estimation results lead us to draw three main conclusions. First, estimates from an appropriately specified and estimated gravity equation *do not* support the notion of a euro effect on trade flows among early euro adopters up to the year 2002. Earlier estimates appear to have been upward-biased largely because they were derived from log-linearised gravity equations estimated by OLS. Second, the finding of “no euro effect” is robust to an extension of the samples used in earlier studies in order to incorporate the most recent data (up to 2013), covering the experience of five subsequent euro joiners. Third, while there is no robust evidence of a euro effect, there is evidence that intra-EU trade flows have expanded faster than the global average during the 2002-2013 period.

Given these findings, our pseudo out-of-sample forecasts compare trade-flow predictions derived under a no-euro-effect assumption with predictions based on a positive euro effect drawn from the middle of the range of traditional estimates. We show that, for the six most recent euro joiners in our sample, the no-effect forecasts of the change in trade flows after euro accession clearly outperform forecasts based on the expectation of a significantly positive euro effect. Therefore, a careful re-examination of the best-available evidence to date seems to suggest that candidate accession countries should not expect a significant boost to their trade from euro adoption.

Academic interest in the trade effects of currency unions peaked around the time of the introduction of the euro in the late 1990s and early 2000s. In an influential paper published shortly after the birth of the euro, Glick and Rose (2002) use a large panel of bilateral trade flows, covering 217 countries in the period 1948-1997, to estimate the impact of currency unions on the trade flows of their member countries. They find a significant and very large effect of pre-euro currency unions on trade. Micco et al. (2003) were the first to use post-EMU data in order to provide an assessment of the specific trade benefits of euro adoption. Their paper finds a 5-15% increase in trade flows in the wake of euro adoption, much smaller than the effect documented by Glick and Rose (2002) for pre-euro currency unions, but still economically and statistically significant. Subsequent research from the early years of EMU, surveyed in Baldwin (2006), points to a similarly-sized euro effect.⁴ Our paper follows in the footsteps of this literature, focusing mainly on the pro-trade effects of the euro. However, we make use of the availability of more recent data and cast the spotlight on countries which were not among the original adopters of the European common currency.

Re-visiting the euro effect in the light of more recent data seems desirable for two reasons. First, we can better assess the long-term impact of the euro on trade flows among original euro adopters than earlier studies whose data sample did not extend beyond the early 2000s. Second, the

⁴Consistently, a recent study by Eicher and Henn (2011), using a large panel data set for the period 1950-2000 in order to evaluate the pro-trade effects of different currency unions individually, confirms that the euro appears to have increased trade among its members by less than other currency unions.

additional euro accessions covered by our sample provide an ideal testing ground for the robustness and predictive power of earlier estimates of the euro effect. Moreover, their experience should be more representative of future euro adopters which, in contrast to the original member countries, are unlikely to join the single currency *en masse*, and are likely to be economically small relative to the eurozone.⁵

The remainder of the paper is structured as follows. Section 2 offers a brief review of the theory underpinning our estimations, since it is crucial for understanding both our estimation methodology and forecasting exercises. It also describes the data used throughout the paper. Section 3 details our estimation results. Section 4 explains, and presents the results of, our pseudo out-of-sample forecasts. Section 5 concludes.

2 Methodology and Data

2.1 Methodology

Like most early studies on the pro-trade effects of common currencies in general, and the euro in particular, we rely on the workhorse model for the analysis of international trade flows – the gravity equation. Since the appropriate empirical application of this model is at the heart of our paper, we will briefly review the underlying theory in this section.

Tinbergen (1962) first documented that international trade flows display empirical regularities resembling Newton’s universal law of gravitation – that is, the value of trade flows between any pair of countries is positively related to measures of their economic “mass”, and negatively to measures of bilateral trade frictions. Yet, the theoretical underpinnings of this observation have only recently become fully understood. In two seminal papers, Anderson (1979) and Anderson and Wincoop (2003) derive a gravity equation of international trade from an Armington model in which consumers have CES preferences over tradable goods that are differentiated by their country of origin. Under these assumptions, they show that the dollar value of imports by country j from country i in year t , M_{ijt} , obeys the following gravity-like relationship:

$$M_{ijt} = \frac{d_{ijt}^\sigma}{P_{it}P_{jt}} \frac{Y_{it}Y_{jt}}{Y_t} \quad \text{with} \quad P_{it} \equiv \sum_j \frac{d_{ijt}^\sigma}{P_{jt}} \frac{Y_{jt}}{Y_t}, \quad (1)$$

⁵In two recent papers, Glick and Rose (2016) and Glick (2017) provide a re-assessment of their earlier estimates of the effects of currency unions using a panel of bilateral trade flows for over 200 countries which covers the period 1948-2013. They generally find significant positive effects of currency unions on trade which are, however, smaller than those reported in Glick and Rose (2002). Moreover, differentiating between early and late euro joiners, Glick (2017) finds no evidence that the euro has boosted trade between the former and the latter. While their large dataset allows them to compare the pro-trade effects of EMU with other currency unions, it does not permit them to estimate a “state of the art” gravity equation – using both a fully array of fixed effects and PPML estimation – for computational reasons. By contrast, since we are primarily interested in the euro effect in its own right, we restrict our sample to 153 countries in the period 1992-2013 and are able to undertake a “state of the art” gravity estimation for this smaller panel.

where Y_{it} denotes the dollar value of country- i output in year t , and Y_t denotes the dollar value of world output in year t . The term d_{ijt} represents the total ad-valorem value of trade costs between countries i and j in year t , and the parameter $\sigma < 0$ captures the responsiveness of trade flows to changes in ad-valorem trade costs (the “trade elasticity”). Subsequent advances in the theory of international trade have documented that equation (1) can be derived from single-sector variants of a range of widely used quantitative trade models.⁶

Empirical applications of equation (1) commonly assume that d_{ijt} can be written as a log-linear function of observables:⁷

$$d_{ijt} = \theta_{ij} \prod_n (z_{ijt}^n)^{\gamma_n}, \quad (2)$$

where θ_{ij} represents the contribution of time-invariant factors to trade frictions (such as i and j sharing a common border) and $(z_{ijt}^n)^{\gamma_n}$ represents the contribution of time-varying factor z_{ijt}^n (such as i and j sharing a common currency). Substituting (2) into (1), we obtain:

$$M_{ijt} = \theta_{ij}^\sigma \frac{\prod_n (z_{ijt}^n)^{\beta_n}}{P_{it} P_{jt}} \frac{Y_{it} Y_{jt}}{Y_t} \quad \text{with} \quad P_{it} \equiv \sum_j \theta_{ij}^\sigma \frac{\prod_n (z_{ijt}^n)^{\beta_n}}{P_{jt}} \frac{Y_{jt}}{Y_t}, \quad (3)$$

where $\beta_n \equiv \sigma \gamma_n$ captures the effect of time-varying factor z_{ijt}^n on bilateral trade flows between countries i and j . Note that our preoccupation in this paper is the value of a particular β_n : the impact of i and j being euro members on their bilateral trade flows.

A key contribution of Anderson and Wincoop (2003) is to highlight the importance of controlling for the so-called “multilateral resistance terms”, P_{it} and P_{jt} . As P_{it} and P_{jt} are functions of the trade costs of countries i and j with all other countries, any gravity-based estimate of $\{\beta_n\}_n$ which does not control for the multilateral resistance terms will suffer from omitted variable bias. As Baldwin and Taglioni (2007) point out, this constitutes a shortcoming of many early estimates of the euro effect on trade. To avoid this problem, we estimate an empirical model of the form:

$$M_{ijt} = \exp \left\{ \delta_{ij} + \delta_{it} + \delta_{jt} + \sum_n \beta_n \ln z_{ijt}^n \right\} \times \eta_{ijt}, \quad (4)$$

where η_{ijt} is an error term with $E[\eta_{ijt} | \{z_{ijt}^n\}] = 1$. The specification in (4) controls for all unobserved variation in time-invariant trade costs across country pairs through country-pair dummies (δ_{ij}), and for country i ’s and j ’s multilateral resistance terms and outputs – as well as all other country-specific shocks – through country-year dummies (δ_{it} , δ_{jt}). The magnitude of each β_n is thus only identified from within-country-pair, over-time variation in trade flows.

Traditional empirical applications of the gravity equation tended to proceed by log-linearising equation (4) and estimating the parameters of interest by least squares. However, as emphasised by Santos Silva and Tenreyro (2006), such parameter estimates will not be consistent in the likely

⁶See Costinot and Rodríguez-Clare (2014).

⁷See, for example, Harrigan (1993), Hummels (2001), Head and Ries (2001), and Baier and Bergstrand (2001) for papers which make this assumption explicitly.

case that the variance of η_{ijt} is a function of other right-hand-side variables. Since the expected value of the log of a random variable depends both on its mean and the higher-order moments of its distribution, the assumption that $E[\ln \eta_{ijt}]$ is independent of the regressors would be violated in this case. In order to ensure the consistency of our estimates, we follow Santos Silva and Tenreyro (2006) and estimate (4) in levels using the Poisson pseudo-maximum-likelihood (PPML) technique.⁸ This has the additional advantage that we do not need to drop or arbitrarily replace zero-value trade flows in our preferred estimation.

2.2 Data

In order to estimate (4), we require annual data on the dollar value of countries' bilateral trade flows, M_{ijt} , as well as the appropriate set of regressors $\{z_{ijt}^n\}$ capturing country-pair-time-varying trade-cost determinants. Data on the value of bilateral trade in goods for the period 1992-2013 is taken from the IMF's *Direction of Trade Statistics*. We choose 1992 as the first year of our sample for consistency with earlier empirical studies on the euro effect, and cover all subsequent years for which trade data is available to date. We exclusively use importer-reported trade flows which are generally regarded as less prone to measurement error than exporter-reported flows.⁹ Our baseline estimates are obtained from a gravity equation which, following Anderson and Wincoop (2003), imposes assumptions that ensure balanced trade between country pairs. Since bilateral trade flows are not balanced in the data, we artificially balance them by imposing $M_{ijt} = M_{jit} = imports_{ijt}^{\frac{1}{2}} \times imports_{jit}^{\frac{1}{2}}$, where $imports_{ijt}$ denotes the dollar value of imports by country j from country i in year t , as reported by j . We drop duplicate observations as well as all observations for which either $imports_{ijt}$ or $imports_{jit}$ is missing.¹⁰

Estimating a balanced-trade gravity equation in this way is consistent with the original literature on the euro effect and significantly reduces the computational demands of the PPML estimations. Yet the recent literature has emphasised that it is preferable to dispense with the assumption of balanced trade in estimating gravity equations, and to use one-way trade flows as the dependent variable (see Head and Mayer, 2014). In the appendix, we re-estimate our main regressions using an unbalanced gravity equation with $M_{ijt} = imports_{ijt}$ as the dependent variable. As shown there, our findings are robust to this alternative estimation approach.

Our main regressor of interest, $\ln z_{ijt}^{EURO}$, is a dummy variable taking value 1 if countries i and j are both euro members in year t , and 0 otherwise. Our main control variable, $\ln z_{ijt}^{EU}$, is a dummy variable taking value 1 if countries i and j are both EU members in year t , and 0 otherwise. We also control for a potential linear EU-specific trend in trade flows, the presence of a non-euro

⁸Due to the large number of country-pair dummies, we do not employ the `ppml` command for STATA provided by Santos Silva and Tenreyro (2006) but the `xtpoisson` command with robust standard errors available in STATA 14.

⁹However, our results are unaffected if we rely wholly or partially on exporter-reported trade flows instead.

¹⁰Given our definition of the dependent variable, the maximum number of observations of trade flows between country pairs in a sample of I countries covering T years is $I(I-1)T/2$.

currency union between i and j , and the presence of a non-EU free trade agreement between i and j . Data on regional trade agreements and common currency areas comes from the WTO and *The Statesman's Yearbook*.¹¹ We also use data on countries' GDPs and population sizes which we obtain from the World Bank's *World Development Indicators*.

To limit the size of the sample, and computational power required to estimate the large number of fixed effects in equation (4), we drop all countries with a population of less than 250,000 inhabitants. We also drop country pairs for which data on trade flows, GDP and our regressors was missing for more than 4 of our 22 sample years. The resulting data set still covers 153 countries and a total of 9,408 country pairs, including bilateral trade between all EU countries as of the end of our sample period.¹²

3 Results

3.1 Overview

In this section, we first provide an estimate of the euro effect on trade using data for the period 1992-2002. Restricting our sample in this way allows us to compare our estimate with the findings from earlier studies which only use data up to 2002. Of course, any euro effect found in this short sample will solely reflect the early experience of the original 11 eurozone members, which formally adopted the common currency in 1999, and Greece, which followed in 2001. We show that, once estimated in line with the methodological discussion in Section 2.1, there is *no* evidence of a positive effect of euro adoption on trade between these countries.

We then extend the sample period up to 2013, now covering a longer time span for the original euro adopters and the euro accession of 5 new member countries which joined between 2002 and 2013. As in the shorter sample period, we find that the hypothesis of “no euro effect” cannot be rejected at reasonable levels of statistical significance. This finding is robust to alternative specifications of the euro dummy and the EU-specific trend which we include in our controls. Indeed, while we find no support for a euro effect of trade in the data, we do find evidence that intra-EU trade flows have expanded faster than the global average during the 2002-2013 period.

¹¹We used data and code kindly provided by Jose de Sousa, and described in De Sousa (2012). The bilateral trade-agreement dummy is based on regional trade areas recognised by the WTO. The bilateral common-currency dummy follows the definition of Glick and Rose (2002) and covers unilateral “dollarisations”, multilateral currency unions and prolonged periods during which the currencies of two countries could be exchanged at a 1:1 par.

¹²Note that we aggregate Belgium and Luxembourg into a single country for the entire 1992-2013 period because pre-1997 trade flows are only reported jointly for “Belgium-Luxembourg” in the *Direction of Trade Statistics*.

3.2 1992-2002 Sample

3.2.1 Estimation Results

To illustrate how and why our estimate of the euro effect differs from the findings of earlier studies, we begin by estimating the impact of euro adoption on trade flows using a specification in the spirit of some of those earlier papers. We regress the log of bilateral trade flows on a euro dummy and a number of control variables using OLS and employing country-pair and year dummies, but *not* country-year dummies. We also use a limited sample of countries, covering only EU members (as of 2002) and eight developed economies – Australia, Canada, Iceland, Japan, New Zealand, Norway, Switzerland and the US. The results are shown in column 1 of Table 1.¹³

[Insert Table 1 about here]

As can be seen from the table, we obtain a roughly 9% increase in the bilateral trade flows among eurozone countries after euro adoption. This is very much in line with a survey of early estimates made by Baldwin (2006) of the pro-trade effects of the euro, which finds it to be “somewhere between 5% and 15%, with 9% being the best estimate.” We now proceed by extending the sample and refining the empirical estimation strategy until we reach our preferred specification. Our discussion will focus on the estimated euro effect, but we will comment on some of the other coefficient estimates below.

We first expand the country coverage of our sample to include all countries which satisfy the criteria described in Section 2.2. Since the extended country sample covers many smaller developing countries, it contains a large number of country pairs and years for which the value of trade flows is reported to be zero (roughly 25% of all observations). For the purpose of the log-linear least-squares estimations, we drop all observations of zero-valued trade flows – but we subsequently re-introduce them in our PPML estimations.

The results from the estimation with expanded country coverage are reported in column 2 of Table 1. Despite the increase in the number of observations – from 2,537 to 62,621 – the coefficient on the euro dummy does not change much: if anything, we find a somewhat larger euro effect at 12%. This is unsurprising as the vast majority of countries now added to the sample are not EU members, and none of them are euro members as of 2002. In column 3, we add two more dummy variables to control for the presence of a (non-EU) free trade agreement and a (non-euro) currency union. Doing so reduces the magnitude of the euro effect back to about 8%.

We now introduce country-year fixed effects in column 4, to account for the multilateral resistance terms which were shown to be an important feature of theory-consistent gravity equations in

¹³The specification and sample underlying the results reported in column 1 of Table 1 correspond exactly to the specification and sample underlying column 2 of Table 1a in the July 2003 IAD Working Paper version of Micco et al. (2003). A comparison of estimates shows that we replicate their findings very closely. We use Micco et al. (2003) as representative of some of the early attempts to assess the euro’s effect on trade because their paper is widely cited for undertaking the first formal empirical assessment of the size of the euro effect specifically (as opposed to the effect of common currencies more generally).

Section 2.1. This once again raises the estimate of the euro effect, to about 14%. Up to this point, we have thus found a range of estimates of the euro effect, but all well within the 5-15% ballpark described by Baldwin (2006).

In columns 5 and 6 we finally move on to our preferred specification, estimating the gravity equation with theory-consistent fixed effects *and* in levels using the PPML technique. PPML estimation of the gravity equation resolves possible coefficient biases in the presence of heteroskedastic standard errors pointed out by Santos Silva and Tenreyro (2006) and allows us to re-introduce observations with zero-valued trade flows back into our sample. To assess which of these has a bigger impact on our estimates, we first provide PPML estimates of our coefficients with zero-valued observations excluded in column 5, and then report PPML estimates based on all available observations in column 6.

As shown in column 5, the magnitude of the estimated euro effect is now reduced dramatically, and it is no longer possible to reject the null hypothesis of “no euro effect” at conventional significance levels. Moreover, none of the coefficient estimates are fundamentally altered when we include zero-valued trade-flow observations in our estimation in column 6. This suggests that the key difference in results between the log-linear estimations and our preferred PPML specification is that the latter addresses a coefficient bias resulting from a correlation between our regressors and higher-order moments of the error term. Below, we provide some evidence that the existence of such a correlation is a legitimate concern when estimating the euro effect. This leads us to our first main conclusion: *estimates from an appropriately specified and estimated gravity equation do not support the notion of a euro effect on trade flows based on the experience of early euro adopters up to the year 2002.*¹⁴

There are a few more noteworthy findings from our estimations for the 1992-2002 period. First, we find a robust economically and statistically significant effect of (non-EU) trade agreements on bilateral trade flows. Second, our preferred specification points towards a one-off 25-28% increase in bilateral trade flows as a result of EU membership, but the evidence of an EU-specific trend in trade flows is mixed at best. Third, and most intriguingly, we cannot reject the hypothesis that the effect of non-euro currency unions on trade is also zero, and that it is equal to the euro effect. This suggests that the absence of a euro effect on trade flows may not be an unusual feature of the euro during the 1992-2002 period.¹⁵

3.2.2 Log-Linear versus PPML Estimation: Evidence of Heteroskedasticity

To understand the difference between the euro-effect estimates derived from log-linear and PPML estimations note, once again, that these do *not* arise because the latter method allows us to

¹⁴Santos Silva and Tenreyro (2010) reach a similar conclusion based on a PPML estimation of the gravity equation for the period 1993-2007. However, they use a much smaller sample of countries and do not employ country-pair fixed effects, which make their estimates of the effect of various trade-cost determinants more vulnerable to the charge of omitted variable bias.

¹⁵This is consistent with the findings of De Sousa (2012) who documents that the currency union effect on trade appears to have declined over time, and had fallen to (or even below) zero by the late 1990s and early 2000s.

incorporate zero-valued trade flows into the sample. Instead, log-linear and PPML estimations yield different estimates of the euro effect even when performed on the *same* sample of trade flows with the *same* control variables (column 4 versus column 5 in Table 1). Santos Silva and Tenreyro (2006) argue that this may be the case when the error term of the gravity equation is heteroskedastic, causing biases in the log-linear estimation of coefficients of interest due to a correlation between the logged error and the regressors. Prompted by our findings, we investigate whether heteroskedasticity is likely to be a concern when estimating the impact of euro adoption on trade flows using a large sample of countries.

[Insert Figure 2 about here]

Figure 2 plots the residual from the estimation in column 5 (expressed as a percentage of predicted trade flows) against the predicted dollar value of bilateral trade flows.¹⁶ The figure strongly indicates systematic differences in the variance of the residual across country pairs: the prediction errors appear to be relatively larger for country pairs which trade little with each other. A more formal test confirms this observation. Regressing the absolute value of the percentage residual on the predicted dollar value of bilateral trade flows, we find that a reduction in predicted trade flows by 1bn US\$ is associated on average with a 1.2 percentage point increase in the absolute percentage residual. This association is statistically significant at the 1% level. It suggests that coefficient estimates obtained by PPML may in general be more reliable when estimating a gravity equation on a large sample of heterogeneous country pairs.

[Insert Figure 3 about here]

In order to investigate whether log-linear estimates of the euro effect specifically could be vulnerable to biases arising from heteroskedasticity, we also check whether euro membership is associated with a systematically different variance of the error term. Figure 3 provides box plots of percentage residuals for two groups of observations: country-pair-year observations of trade flows for country pairs in the EU but not in the Eurozone (left-hand side), and for country pairs in the EU *and* in the Eurozone (right-hand side). The median residual is very close to zero for both groups, but the figure shows that the variance of percentage residuals is noticeably larger for EU-and-Eurozone observations than for EU-observations only. Again, a formal test confirms this: regressing the absolute percentage residual for the 892 intra-EU country-pair-year observations on the euro dummy, we obtain a coefficient of 1.31, which is statistically significant at the 5% level.

This evidence causes us to place greater faith in the heteroskedasticity-robust estimates of the euro effect obtained from PPML estimations of the gravity equation, underpinning our conclusion that there is no empirical support for a significant euro effect on trade flows up to the year 2002. In the next subsection, we show that the finding of “no euro effect” is robust to an extension of the sample period which allows us to account for the experience of more recent euro adopters. In section

¹⁶Note that, for expositional purposes only, the figure omits the largest and smallest 5% of residuals.

3.4, we show that it is robust to different compositions of the country sample. Throughout, the PPML-based estimates deliver the most consistent verdict of the size of the euro effect, while OLS-based estimates vary widely. This is a further reason PPML constitutes our preferred estimation method: unlike with OLS, the magnitude and statistical significance of the euro effect obtained from a theory-consistent gravity equation estimated with PPML does not appear to be sensitive to the precise sample composition.

3.3 1992-2013 Sample

Table 2 presents the results from repeating the same six estimations reported in Table 1 after extending our sample to the full 1992-2013 period – which raises the maximum number of observations from 83,466 to 184,642, adds an additional 11 years of euro membership for the early adopters and covers 5 additional euro accessions. The table serves to illustrate two points.

[Insert Table 2 about here]

First, extending the sample period leads to very different conclusions about the magnitude of the euro effect derived from log-linear specifications in columns 1-4. This change is especially dramatic for specifications employing the full country-pair sample, and country-time fixed effects: the euro effect is now found to be negative and significant at -8% in column 3, and negative and insignificant in column 4.

Second, the finding of “no euro effect” from the PPML estimations in columns 5 and 6 remains robust. The euro coefficient is now also negative in these specifications but remains relatively small in absolute magnitude and statistically insignificant at conventional levels. Out of the theory-consistent estimates of the euro on trade using gravity equations, the PPML estimations thus appear to provide the most robust assessment of the magnitude of the effect across sample periods. This gives rise to the second main conclusion of our analysis: *estimates from an appropriately specified and estimated gravity equation do not support the notion of a euro effect on trade flows based on the experience of euro adopters to date.*

A major difference between our 1992-2002 and 1992-2013 PPML-based estimates is that we find a positive and statistically significant EU-specific trend in bilateral trade flows in the longer period. We explore the robustness of this finding in Section 3.4.3. The estimated pro-trade effect of (non-EU) free trade agreements is also significantly reduced in the longer period, but our estimate of the trade benefit of (non-euro) currency unions is virtually unchanged in magnitude and statistical significance, and we still cannot reject the hypothesis that it is equal to the euro effect, and equal to zero. We present a number of robustness checks in Section 3.4.

3.4 Robustness

3.4.1 Country Sample

In a recent survey of estimates of the EMU effect on trade, Rose (2016) finds that the size of the euro effect estimated by OLS from a theory-consistent log-linearised gravity equation appears to be sensitive to the choice of country sample. In particular, he shows that the euro effect is the largest when a comprehensive country sample is used, and that the magnitude of the estimate declines for samples comprised of only high-income or EU countries.

We proceed to check whether our main finding of “no euro effect” is robust to variations in the country sample and report the results in Table 3. Following Rose (2016), we re-estimate our gravity equation for two subsamples of our main dataset: one containing only high-income countries (using the current World Bank definition of countries with a real per-capita GDP in excess of \$12,736); and one containing only countries which are EU members for at least one year between 1992 and 2013. To highlight the differences between log-linear OLS and our preferred PPML estimation, we present results from the former (in columns 1 and 4) alongside results from the latter, with zero-valued trade flows omitted (in columns 2 and 5) and included (in columns 3 and 6). Throughout, we use the full sample period 1992-2013.

[Insert Table 3 about here]

As Table 3 shows, the hypothesis of a zero euro effect cannot be rejected in *any* of these samples, whether estimated by OLS or PPML. It also confirms the tendency highlighted in Rose (2016): among OLS estimates, the smallest euro effect (-9%) is found when we use the smallest country sample (EU members only, in column 4). However, this tendency is not evident in our PPML estimates. Rose (2016) suggests that the OLS-estimated euro effect declines with the country coverage of the sample because omitting smaller and poorer countries results in biased estimates of the multilateral resistance terms. Our findings indicate an alternative explanation: since the variance of the residual is relatively larger for country pairs whose predicted trade flows are smaller in absolute magnitude (as seen in Figure 2), and since those pairs tend to be comprised of smaller and poorer countries, dropping those countries alters the OLS estimate of the euro effect by reducing the extent of heteroskedasticity problem.

[Insert Figure 4 about here]

Figure 4 provides a graphical overview of the different point estimates and 95%-confidence intervals of the euro effect obtained from theory-consistent gravity equations using the OLS and PPML estimators on different subsamples of our full data set.¹⁷ The figure highlights the robustness of the “no euro effect” finding. It also shows that, across the board, our preferred PPML

¹⁷The left-hand panel is based on the estimates of the euro effect reported in column 4 of Table 2, column 4 of Table 1, column 1 of Table 3 and column 4 of Table 3, respectively. The right-hand panel is based on column 5 of Table 2, column 5 of Table 1 column 2 of Table 3 and column 5 of Table 3, respectively.

estimation method results in estimates which are smaller in absolute magnitude, more precise and less variable.¹⁸

3.4.2 Heterogeneous Euro Effect

Our findings raise the possibility that the EMU effect on the trade flows of the euro’s early adopters may have been different from its impact on later joiners. In columns 3-5 of Table 4, we check for heterogeneity in the euro’s effect on trade across periods and country pairs. In column 3, we add a second euro dummy which takes the value 1 if countries i and j are both euro members and the year is greater than 2002, and 0 otherwise. The purpose is to assess whether the euro effect appears to have become stronger or weaker in the post-2002 portion of our sample. In column 4, we add an additional euro dummy which takes the value 1 if countries i and j are both euro members and at least one of them joined the euro since 2002, and 0 otherwise. This is to investigate whether the euro effect appears to have been different for late joiners. In column 5, both of these dummies are jointly added to the estimation.

[Insert Table 4 about here]

While the inclusion of the post-2002 euro dummy reduces the coefficient on the main euro dummy away from zero, making it marginally statistically significant, the additional euro dummies are not statistically significant in columns 3-5, and an F -test comfortably rejects the joint significance of the euro dummies in all three cases. Therefore, allowing for heterogeneous euro effects does little to alter our conclusion that there is no evidence of a euro effect on trade flows.

3.4.3 EU Trend

So far, we have allowed for a linear EU-specific trend in bilateral trade flows. The estimation results presented in Tables 1, 2 and 3 highlight that the finding of an EU trend appears to be highly sensitive to the choice of estimation method and sample composition. In the following, we explore alternative specifications of the EU effect on trade, in order to investigate if the precise specification affects our estimate of the euro effect.

[Insert Table 5 about here]

Table 5 reports results for four alternative theory-consistent PPML estimations for the full sample period – one with an EU dummy only (column 1), one with a linear EU trend (column 2), one with a quadratic EU trend (column 3) and one with EU-year dummies (column 4). As the table shows, the specification of the EU trend does have an impact on the estimated euro effect: without allowing for an EU trend, the euro effect is estimated to be about 7%, and statistically significant

¹⁸In a recent manuscript, Larch et al. (2017) document that this finding also extends to a sample of 200 countries over the 1948-2013 period, using a new PPML estimator which allows them to estimate a gravity equation with theory-consistent fixed effects for a much larger data set.

at the 10% level. For any other specification of the EU trend, however, the euro coefficient is much closer to zero and the euro effect is not statistically significant.

[Insert Figure 5 about here]

The four different EU trends implied by the four different specifications are plotted in Figure 5. The figure highlights two features of the behaviour of EU trade flows over the 1992-2013 period. First, the linear, quadratic and dummy specifications all point towards a large increase in intra-EU trade flows of 27-30% relative to the global average between 1992 and 2013.¹⁹ Second, the quadratic and dummy specifications suggest that this increase only manifested itself in the latter third of this period, with no – or negative – EU-specific trade growth between 1992 and the early 2000s. This explains why allowing for a linear EU trend resulted in the finding of a weakly negative trend in our preferred specifications in Table 1, but a strongly positive trend in columns 2-5 of Table 2. On this basis, we reach the third and final conclusion from our estimations: *there is evidence that intra-EU trade flows have expanded faster than the global average during the 2002-2013 period.*²⁰

[Insert Figure 6 about here]

This final finding raises the question whether the absence of a euro effect in our preferred specifications is due to the use of the PPML estimator or the inclusion of an EU trend. We attempt to address this question by re-estimating the euro effects from different samples and estimation methods reported in Figure 4, this time including only an EU dummy – but *not* an EU trend – among the control variables. Figure 6 provides a graphical overview of the updated estimates and 95%-confidence intervals. As before, the euro-effect estimates obtained from the log-linear estimations vary substantially, and we obtain positive and significant euro effects in the full and 1992-2002 samples, but also a significantly *negative* effect in the EU-only sample. By contrast, the PPML estimations deliver estimates which are consistently close to zero, and never statistically significant at the 5% level. Therefore, while there is evidence of a positive EU trend in the 1992-2013 period, our finding of “no euro effect” using the PPML estimator prevails even if we do not control for this trend.

4 Pseudo Out-of-Sample Forecasts

4.1 Forecasting Changes in Trade Flows After Eurozone Accession

Earlier estimates of the euro’s effect on member countries’ trade, based mostly on pre-2002 data, found that the euro raised trade flows by 5-15% on average. We showed in the previous section that

¹⁹Note that any trend in average global trade flows should be captured by the array of country-year fixed effects employed in our preferred specification.

²⁰Bun and Klaassen (2007) and Berger and Nitsch (2008) also find that controlling for long-run trends in EU trade flows significantly reduces the estimated magnitude of the euro effect. However, since both papers focus on the pre-2003 period, they only reflect the experience of the early euro adopters. Our results highlight that their finding remains relevant for more recent data, and subsequent eurozone accessions.

this appears to have been an over-estimate, and that an appropriately estimated gravity equation provides no support for a euro effect on trade flows – neither in the pre-2002 period nor based on the latest available data. In this section we return to a question posed in the introduction: were earlier estimates of the euro effect on trade flows a good guide to the euro’s impact on trade for later additions to the eurozone? We attempt to provide an answer by comparing forecasts of the rise in new members’ trade flows using different estimates of the euro effect with their actual trade growth following their euro accession.

To this end, we turn once more to the gravity equation introduced in Section 2.1. Quantitative trade models of the type which are used to derive structural gravity equations along the lines of (1) are now widely used to analyse the welfare effects of trade-policy changes: the estimated effect of a policy change (e.g. euro membership) on trade barriers is fed back into the model in order to make predictions about the general-equilibrium changes in trade and incomes. As a “reality check” of any estimated policy impact on trade, and to bolster confidence in the resulting welfare calculations, it seems desirable to establish that the estimate does not imply changes in the value of trade flows relative to outputs which are systematically at odds with the data. In the following, we show how to assess this for different estimates of the euro effect, relying on the formal structure of the gravity equation in (1).

Note from (1) that

$$M_{iit} = \frac{d_{iit}^\sigma Y_{it}^2}{P_{it}^2 Y_t}, \quad (5)$$

where M_{iit} denotes the dollar value of country i ’s trade with itself in year t . We can now use (5) to eliminate the multilateral resistance terms from equation (1) and obtain

$$\frac{M_{ijt}}{M_{iit}^{\frac{1}{2}} M_{jjt}^{\frac{1}{2}}} = \left(\frac{d_{ijt}}{d_{iit}^{\frac{1}{2}} d_{jjt}^{\frac{1}{2}}} \right)^\sigma. \quad (6)$$

Equation (6) has an intuitive interpretation: the value of trade flows between countries i and j relative to the geometric mean of their “internal” trade in year t depends on the value of trade costs between i and j relative to the geometric mean of their “internal” trade costs. Using the log-linear trade-cost function from (2), and imposing the common assumption that “internal” trade costs are constant,

$$\sum_{t \in T_1} \frac{1}{T_1} \frac{M_{ijt}}{M_{iit}^{\frac{1}{2}} M_{jjt}^{\frac{1}{2}}} \bigg/ \sum_{t \in T_0} \frac{1}{T_0} \frac{M_{ijt}}{M_{iit}^{\frac{1}{2}} M_{jjt}^{\frac{1}{2}}} - 1 = \frac{\sum_{t \in T_1} \Pi_n (z_{ijt}^n)^{\beta_n} / T_1}{\sum_{t \in T_0} \Pi_n (z_{ijt}^n)^{\beta_n} / T_0} - 1 \quad (7)$$

Equation (7) states that the growth in country i ’s and j ’s bilateral trade (relative to their “internal” trade) between the average year in period 0 and period 1 should be driven by changes in the trade cost determinants between these periods. Abusing notation, we let T_0 and T_1 denote the number and set of years in periods 0 and 1, respectively. For arbitrary periods $\{T_0, T_1\}$ the actual change in trade flows – the left-hand side of equation (7) – can be taken directly from the data. The only

additional information required concerns the dollar value of “internal trade” for country i in each year t , M_{iit} . We obtain it by subtracting the sum of the dollar value of country i ’s international trade flows, $\sum_j M_{ijt}$, from the dollar value of country i ’s gross output in year t . The latter is taken from the World Input Output Database (see Timmer et al., 2015).

Our forecast of the change in trade flows between countries i and j after country i ’s euro accession (relative to their “internal” trade) amounts to calculating a value for the right-hand side of (7) for suitably chosen T_0 , T_1 . We determine the predicted post-accession change in trade flows for the six eurozone economies which joined the euro after 1999 and before 2013 – Greece, Slovenia, Cyprus, Malta, the Slovak Republic and Estonia. In the interest of expositional clarity, we only forecast the change in their trade flows with those eurozone economies which adopted the single currency in 1999 – Austria, Belgium-Luxembourg, Finland, France, Germany, Ireland, Italy, the Netherlands, Portugal and Spain. For each of the six “late joiners” we define T_0 as the years from 1996 until the year prior to their euro accession, and T_1 as all years between their euro accession and 2013.

To calculate the forecasted change in trade flows, we need a set of trade-cost determinants $\{z_{ijt}^n\}_n$ and corresponding coefficients $\{\beta_n\}_n$. In addition to the euro dummy, we choose the same trade cost determinants as in our preferred specifications in column 6 of Tables 1 and 2: an EU dummy, a linear EU trend, an non-EU trade treaty dummy, and a non-euro common currency dummy, all as defined in Section 2.2. The corresponding coefficients for each of the six “late joiner” economies are taken from an estimation of the gravity equation, using our preferred specification and the full country sample, up to the year prior to the respective country’s euro entry. Only the euro effect, β_{EURO} , is simply imposed by us: we provide one forecast in which $\beta_{EURO} = 0$ in line with our findings above, and one in which $\beta_{EURO} = .09$ roughly in line with Baldwin’s (2006) “best estimate” in a survey of early studies of the euro effect.

Our forecasts are designed to emulate a particular thought experiment. Suppose each of the “late joiner” economies had attempted to forecast the change in its trade with the original eurozone economies in the year prior to its accession under a particular expectation of the magnitude of the euro effect (and with all other drivers of trade flows estimated from a “state of the art” gravity equation up to that point). We seek to assess whether, in this hypothetical scenario, the expectation of no euro effect would have resulted in a more accurate forecast than the expectation of a 9% euro effect.

4.2 Results

For each of the six recent euro adopters, Figure 7 shows the average change in trade flows, relative to internal trade, with the ten early euro adopters – as seen in the data (black bar), as forecast under the assumption of no euro effect (blue bar), and as forecast under the assumption of a 9% euro effect (red bar).

[Insert Figure 7 about here]

It is easy to see from the figure that the expectation of a 9% euro effect results in a forecast of the average change in trade flows which systematically exceeds the realised average change in the six cases. Moreover, with the exception of Greece, the expectation of no euro effect results in a forecast of the average change which is more closely aligned with the data. Turning to individual bilateral trade flows, Table 6 reports the mean squared error of the two forecasts across changes in trade flows between each late euro adopter and the core eurozone. The mean squared error of the forecast based on the expectation of a 9% euro effect exceeds that of the no-effect forecast for all “recent joiners” except Cyprus (for which the former performs marginally better than the latter). Overall, therefore, the no-effect forecast appears to outperform forecasts based on the expectation of a significantly positive euro effect.

[Insert Table 6 about here]

It should be noted that our comparison of forecasts relies heavily on the assumptions which underpin the standard structural gravity equation described in Section 2.1. A violation of these assumptions would thus undermine our assessment of the relative forecast accuracy of different estimates of the euro effect. For example, the gravity equation we estimate does not allow for any dynamic effects of euro adoption, whereby the full size of the effect may take several years to materialise. It also does not allow for the possibility that trade flows may depart significantly from a standard gravity benchmark during major economic shocks – a particular concern in our case, as most of our “late joiners” became EMU members during a period marked by the global financial crisis and Eurozone debt crisis.

For these reasons, the forecast comparison should not be viewed as definitive evidence of the absence of a euro effect. Nevertheless, it provides a further cautionary signal which reinforces the message from our estimation results from Section 3: given the best-available evidence to date, evaluated in accordance with the implications of standard quantitative trade models, candidate countries for euro accession should not expect euro membership to result in a significant boost to their trade.

5 Conclusion

In this paper, we set out to re-evaluate the trade benefits of euro membership using the latest available data in order to take account of the experience of recent euro adopters. Our empirical analysis yields several striking findings. First, a re-examination of the empirical evidence up to the year 2002 – using “state of the art” gravity methodology – finds no robust support for a euro effect on trade flows, suggesting that the magnitude of the euro effect reported in earlier comparable studies was an over-estimate. Second, the absence of a euro effect of trade flows also appears to be the only robust conclusion which can be drawn from an extended sample, covering 11 additional years of data and five more recent euro accessions. Finally, while evidence of a euro

effect is elusive, we do detect an EU-specific trend in international trade flows. A pseudo out-of-sample forecasting exercise reinforces the message arising from our estimations: the expectation of a significantly positive euro effect would have caused recent euro joiners to over-predict the change in their trade flows after EMU accession.

The results presented here are of relevance to current and possible future accession candidates for EMU membership. There may be a host of political and economic reasons why any one of these countries may wish to join the euro. However, based on the best evidence available to date, we do not believe that increased trade integration with the eurozone should be considered one of them.

In addition, our paper adds to a growing body of evidence suggesting that the trade gains from EMU have been small compared to other currency unions and, more generally, that the monetary unions tend to have heterogeneous effects on trade flows.²¹ Investigating the sources of this heterogeneity theoretically and empirically poses an interesting challenge for research in international macroeconomics going forward. “Trade gains” form a key part of the theory of Optimum Currency Areas (Mundell, 1961). Thus, a better understanding of when these gains are likely to be large would help flesh out an old theory – with practical benefits well beyond the borders of the eurozone.

²¹Recently, Eicher and Henn (2011), De Sousa (2012) and Glick and Rose (2016) all show that the trade-promoting effects of currency unions have varied significantly across space and time.

References

- Anderson, J. E. (1979). A theoretical foundation for the gravity equation. *American Economic Review*, 69(1):106–116.
- Anderson, J. E. and Wincoop, E. V. (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, 93(1):170–192.
- Baier, S. L. and Bergstrand, J. H. (2001). The growth of world trade: Tariffs, transport costs, and income similarity. *Journal of International Economics*, 53(1):1–27.
- Baldwin, R. (2006). In or out: Does it make a difference: An evidence-based analysis of the trade effects of the euro. *Center For Economic Policy Research*.
- Baldwin, R. and Taglioni, D. (2007). Trade effects of the euro: a comparison of estimators. *Journal of Economic Integration*, 22:780–818.
- Berger, H. and Nitsch, V. (2008). Zooming out: The trade effect of the euro in historical perspective. *Journal of International Money and Finance*, 27(8):1244–1260.
- Bun, M. and Klaassen, F. (2007). The euro effect on trade is not as large as commonly thought. *Oxford Bulletin of Economics and Statistics*, 69(4):473–496.
- Costinot, A. and Rodríguez-Clare, A. (2014). Trade theory with numbers: Quantifying the consequences of globalization. *Handbook of International Economics*, ed. by G. Gopinath, E. Helpman, and K. Rogoff, vol. 4, chap. 4., pages 106–116.
- De Sousa, J. (2012). The currency union effect on trade is decreasing over time. *Economics Letters*, 117(3):917–920.
- Eicher, T. S. and Henn, C. (2011). One money, one market: A revised benchmark. *Review of International Economics*, 19(3):419–435.
- Glick, R. (2017). Currency unions and regional trade agreements: Emu and eu effects on trade. *Comparative Economic Studies*, 59(2):194–209.
- Glick, R. and Rose, A. K. (2002). Does a currency union affect trade? the time-series evidence. *European Economic Review*, 46(6):1125–1151.
- Glick, R. and Rose, A. K. (2016). Currency unions and trade: A reassessment. *European Economic Review*, 87(C):78–91.
- Harrigan, J. (1993). Oecd imports and trade barriers in 1983. *Journal of International Economics*, 35(1):91–111.

- Head, K. and Mayer, T. (2014). Gravity equations: Workhorse, toolkit, and cookbook. In Gopinath, G., Helpman, E., and Rogoff, K., editors, *Handbook of International Economics*, volume 4, chapter 3, pages 131–195. Elsevier.
- Head, K. and Ries, J. (2001). Increasing returns versus national product differentiation as an explanation for the pattern of us-canada trade. *American Economic Review*, 9(4):858–876.
- Hummels, D. (2001). Toward a geography of trade costs. *mimeo, Purdue University*.
- Larch, M., Wanner, J., Yotov, Y. V., and Zylkin, T. (2017). The currency union effect: A ppml re-assessment with high-dimensional fixed effects. CESifo Working Paper No. 6464.
- Micco, A., Stein, E., and Ordóñez, G. (2003). The currency union effect on trade: Early evidence from emu. *Economic Policy*, 18(37):315–356.
- Mundell, R. A. (1961). A theory of optimum currency areas. *American Economic Review*, 51(4):657–665.
- Rose, A. K. (2016). Why do estimates of the emu effect on trade vary so much? *Open Economies Review*, forthcoming.
- Santos Silva, J. and Tenreyro, S. (2006). The log of gravity. *The Review of Economics and Statistics*, 88(4):641–658.
- Santos Silva, J. and Tenreyro, S. (2010). Currency unions in prospect and retrospect. *Annual Review of Economics*, 2:51–74.
- Timmer, M. P., Dietzenbacher, E., Los, B., Stehrer, R., and Vries, G. J. (2015). An illustrated user guide to the world input-output database: the case of global automotive production. *Review of International Economics*, 23(3):575–605.
- Tinbergen, J. (1962). Shaping the world economy: Suggestions for an international economic policy. *New York, The Twentieth Century Fund*.

A Appendix: Unbalanced Gravity

A.1 Methodology and Data

Following Anderson and Wincoop (2003), the gravity equation in (3) is derived under assumptions which ensure that bilateral trade flows are balanced: $M_{ijt} = M_{jit}$. Since trade imbalances between country pairs are observed in the data, we can generalise (3) to allow for such imbalances (see Head and Mayer, 2014). The resulting generalisation of (3) takes the following form:

$$M_{ijt} = \theta_{ij}^\sigma \frac{\Pi_n (z_{ijt}^n)^{\beta_n}}{O_{it} P_{jt}} \frac{Y_{it} E_{jt}}{Y_t} \quad \text{with } O_{it} \equiv \sum_j \theta_{ij}^\sigma \frac{\Pi_n (z_{ijt}^n)^{\beta_n}}{P_{jt}^{-\theta}} \frac{E_{jt}}{Y_t}, \quad P_{jt} \equiv \sum_i \theta_{ij}^\sigma \frac{\Pi_n (z_{ijt}^n)^{\beta_n}}{O_{it}} \frac{Y_{it}}{Y_t}, \quad (8)$$

where E_{it} now denotes country- i expenditure (which need not equal country- i output, Y_{it}). Note that this implies that M_{ijt} does not necessarily equal M_{jit} . In order to estimate the euro effect in a manner which is consistent with (8), we need to update the empirical specification in (4) to


$$M_{ijt} = \exp \left\{ \delta_{ij} + \delta_{it}^X + \delta_{jt}^M + \sum_n \beta_n \ln z_{ijt}^n \right\} \times \eta_{ijt}, \quad (9)$$

where δ_{it}^X and δ_{jt}^M respectively represent country- i -as-exporter and country- j -as-importer fixed effects, and δ_{ij} is no longer restricted to equal δ_{ji} .

The specification in (9) has the advantage that we no longer need to balance trade flows artificially by calculating the geometric mean of bilateral imports and exports. Instead, $M_{ijt} = \text{imports}_{ijt}$ is the new dependent variable. On the flip side, estimating (9) is considerably more computationally demanding. It more than doubles the number of observations, since we no longer need to drop observations in which one of the two flows is missing. It also doubles the number of fixed effects which need to be estimated for the empirical gravity model to be consistent with (8).


A.2 Results

We re-estimate all regressions whose results were reported in Tables 1 and 2 using the specification in (9), without artificially balanced trade, instead of (4). The results are reported in Tables A1 and A2, respectively. Across the board, the differences between the coefficient estimates listed in Tables A1 and A2 and those listed in Tables 1 and 2 are negligible. This leads us to conclude that imposing balanced trade, as we do for all the estimations in the main body of the paper, does not appear to bias our findings in a significant or systematic way.



EZ/EU refers to euro/EU members as of 1999. Year prior to euro accession (vertical line) used as base.

Figure 1: Trade with Eurozone/Trade with All EU Countries
for Six Post-1999 Euro Adopters



The scatter plots omits the largest and smallest 5% of residuals.

Figure 2: Gravity-Predicted Trade Flows and Gravity Residuals




Figure 3: Gravity Residuals for Different EU Country Pairs




Figure 4: Euro Effect Estimates from Theory-Consistent Gravity Equations
– Different Samples and Estimation Methods (with EU-Trend)




Figure 5: Alternative EU-Trend Estimates





Figure 6: Euro Effect Estimates from Theory-Consistent Gravity Equations
– Different Samples and Estimation Methods (without EU-Trend)



EZ refers to euro members as of 1999.

Figure 7: Results of Pseudo-Out-of-Sample Forecasts,
Aggregate Trade Growth

Dep. variable: Log bil. trade (PPML: bil. trade)	(1) Sml. Sample, 1992-2002	(2) World Sample, 1992-2002	(3) World Sample, 1992-2002	(4) World Sample, 1992-2002	(5) World Sample, 1992-2002 (PPML: no 0s)	(6) World Sample, 1992-2002 (PPML)
Euro	0.085 (0.026)***	0.122 (0.026)***	0.082 (0.028)***	0.135 (0.053)**	0.018 (0.028)	0.019 (0.029)
Log product of GDPs	2.854 (0.522)***	-0.346 (0.118)***	-0.348 (0.117)***			
Log product of per- capita GDPs	-2.488 (0.539)***	0.835 (0.119)***	0.834 (0.118)***			
EU	-0.032 (0.034)	-0.027 (0.033)	0.183 (0.044)***	0.327 (0.070)***	0.246 (0.050)***	0.280 (0.054)***
EU trend	0.017 (0.005)***	-0.020 (0.004)***	-0.019 (0.004)***	0.009 (0.007)	-0.010 (0.006)*	-0.008 (0.006)
FTA			0.189 (0.029)***	0.175 (0.034)***	0.126 (0.032)***	0.147 (0.033)***
Other com. curr.			-0.124 (0.024)***	-0.027 (0.035)	0.021 (0.017)	0.022 (0.017)
Observations	2,537	61,621	61,621	61,621	61,621	83,466
R^2	1.00	0.94	0.94	0.95	1.00	1.00
Fixed effects:						
- Country-pair	Yes	Yes	Yes	Yes	Yes	Yes
- Country-year	No	No	No	Yes	Yes	Yes
- Time	Yes	Yes	Yes	No	No	No

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$
Note: Standard errors clustered at the country-pair level.

Table 1: Regressions, 1992-2002

Dep. variable: Log bil. trade (PPML: bil. trade)	(1) Sml. Sample, 1992-2013	(2) World Sample, 1992-2013	(3) World Sample, 1992-2013	(4) World Sample, 1992-2013	(5) World Sample, 1992-2013 (PPML: no 0s)	(6) World Sample, 1992-2013 (PPML)
Euro	0.103 (0.034)***	-0.039 (0.038)	-0.081 (0.039)**	-0.062 (0.052)	-0.030 (0.036)	-0.024 (0.037)
Log product of GDPs	0.875 (0.279)***	0.750 (0.051)***	0.769 (0.051)***			
Log product of per-capita GDPs	-0.413 (0.302)	-0.085 (0.052)	-0.102 (0.052)**			
EU	-0.054 (0.041)	0.236 (0.042)***	0.315 (0.045)***	0.331 (0.061)***	0.165 (0.056)***	0.160 (0.056)***
EU trend	0.010 (0.003)***	0.009 (0.003)***	0.011 (0.003)***	0.043 (0.004)***	0.014 (0.003)***	0.014 (0.003)***
FTA			0.148 (0.024)***	0.162 (0.027)***	0.065 (0.026)**	0.063 (0.026)**
Other com. curr.			-0.123 (0.028)***	-0.023 (0.039)	0.022 (0.025)	0.022 (0.025)
Observations	5,078	138,929	138,929	138,929	138,929	184,642
R^2	0.99	0.93	0.93	0.94	1.00	1.00
Fixed effects:						
- Country-pair	Yes	Yes	Yes	Yes	Yes	Yes
- Country-time	No	No	No	Yes	Yes	Yes
- Time	Yes	Yes	Yes	No	No	No

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$
Note: Standard errors clustered at the country-pair level.

Table 2: Regressions, 1992-2013

Dep. variable: Log bil. trade (PPML: bil. trade)	(4) HI only, 1992-2013 (Log-linear)	(5) HI only, 1992-2013 (PPML: no 0s)	(6) HI only, 1992-2013 (PPML)	(4) EU only, 1992-2013 (Log-linear)	(5) EU only, 1992-2013 (PPML: no 0s)	(6) EU only, 1992-2013 (PPML)
Euro	-0.056 (0.056)	-0.046 (0.039)	-0.043 (0.040)	-0.086 (0.055)	0.034 (0.048)	0.046 (0.052)
EU	0.008 (0.081)	0.036 (0.076)	0.041 (0.076)	0.278 (0.122)**	0.662 (0.221)***	0.666 (0.219)***
EU trend	0.023 (0.006)***	0.009 (0.004)**	0.009 (0.004)**	-0.036 (0.013)***	-0.041 (0.015)***	-0.040 (0.014)***
FTA	-0.015 (0.056)	-0.069 (0.039)*	-0.068 (0.039)*	0.104 (0.063)*	0.367 (0.128)***	0.377 (0.129)***
Other com. curr.	-0.030 (0.044)	0.013 (0.026)	0.014 (0.026)	-0.054 (0.044)	-0.012 (0.031)	-0.012 (0.031)
Observations	21,890	21,890	23,634	6,899	6,899	7,037
R^2	0.97	1.00	1.00	0.99	1.00	1.00
Fixed effects:						
- Country-pair	Yes	Yes	Yes	Yes	Yes	Yes
- Country-year	Yes	Yes	Yes	Yes	Yes	Yes

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$
Note: Standard errors clustered at the country-pair level.

Table 3: Robustness Checks – Country Sample

Dep. variable:	(1)	(2)	(3)	(4)	(5)
Bilateral trade	World Sample, 1992-2002 (PPML)	World Sample, 1992-2013 (PPML)	World Sample, 1992-2013 (PPML)	World Sample, 1992-2013 (PPML)	World Sample, 1992-2013 (PPML)
Euro	0.019 (0.029)	-0.024 (0.037)	-0.057 (0.031)*	-0.027 (0.040)	-0.059 (0.033)*
Euro (post 2002)			0.046 (0.029)		0.046 (0.030)
Euro (new members)				0.022 (0.096)	0.011 (0.098)
EU	0.280 (0.054)***	0.160 (0.056)***	0.175 (0.057)***	0.157 (0.054)***	0.173 (0.056)***
EU trend	-0.008 (0.006)	0.014 (0.003)***	0.013 (0.003)***	0.014 (0.003)***	0.013 (0.003)***
FTA	0.147 (0.033)***	0.063 (0.026)**	0.063 (0.026)**	0.063 (0.026)**	0.063 (0.026)**
Other com. curr.	0.022 (0.017)	0.022 (0.025)	0.022 (0.025)	0.022 (0.025)	0.022 (0.025)
Observations	83,466	184,642	184,642	184,642	184,642
R^2	1.00	1.00	1.00	1.00	1.00
Fixed effects:					
- Country-pair	Yes	Yes	Yes	Yes	Yes
- Country-year	Yes	Yes	Yes	Yes	Yes

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$
Note: Standard errors clustered at the country-pair level.

Table 4: Robustness Checks – Heterogeneous Euro Effect

Dep. variable: Bilateral trade	(1) World Sample, 1992-2013 (PPML)	(2) World Sample, 1992-2013 (PPML)	(3) World Sample, 1992-2013 (PPML)	(4) World Sample, 1992-2013 (PPML)
Euro	0.066 (0.034)*	-0.024 (0.037)	0.002 (0.038)	0.007 (0.041)
FTA	0.063 (0.026)**	0.063 (0.026)**	0.062 (0.026)**	0.061 (0.026)**
Other com. curr.	0.024 (0.024)	0.022 (0.025)	0.023 (0.024)	0.023 (0.024)
EU	0.307 (0.045)***	0.160 (0.056)***	0.256 (0.065)***	
EU trend		0.014 (0.003)**	-0.010 (0.008)	
EU trend ²			0.001 (0.000)***	
Observations	184,642	184,642	184,642	184,642
R ²	1.00	1.00	1.00	1.00
Fixed effects:				
- Country-pair	Yes	Yes	Yes	Yes
- Country-year	Yes	Yes	Yes	Yes
- EU-year	No	No	No	Yes

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$
Note: Standard errors clustered at the country-pair level.

Table 5: Robustness Checks – EU Trend

Country	Euro Accession	MSQE (no euro effect)	MSQE (9% euro effect)
Greece	2001	1987.7	2743.1
Slovenia	2007	1526.4	2442.5
Cyprus	2008	178.9	173.0
Malta	2008	1659.5	2406.3
Slovak Republic	2009	855.5	970.9
Estonia	2011	432.5	547.2

Table 6: Results of Pseudo-Out-of-Sample Forecasts, Mean-Squared Errors (MSQE)
for Country-Pair Trade Growth

Dep. variable: Log imports (PPML: imports)	(1) Sml. Sample, 1992-2002	(2) World Sample, 1992-2002	(3) World Sample, 1992-2002	(4) World Sample, 1992-2002	(5) World Sample, 1992-2002 (PPML: no 0s)	(6) World Sample, 1992-2002 (PPML)
Euro	0.084 (0.025)***	0.126 (0.025)***	0.084 (0.027)***	0.153 (0.052)***	0.023 (0.030)	0.023 (0.030)
Log product of GDPs	2.857 (0.518)***	-0.337 (0.108)***	-0.340 (0.108)***			
Log product of per-capita GDPs	-2.491 (0.537)***	0.781 (0.109)***	0.782 (0.109)***			
EU	-0.031 (0.031)	-0.028 (0.031)	0.161 (0.043)***	0.319 (0.068)***	0.220 (0.051)***	0.246 (0.052)***
EU trend	0.018 (0.005)***	-0.022 (0.004)***	-0.021 (0.004)***	0.002 (0.007)	-0.013 (0.005)**	-0.011 (0.005)**
FTA			0.166 (0.031)***	0.154 (0.035)***	0.100 (0.025)***	0.124 (0.027)***
Other com. curr.			-0.130 (0.025)***	-0.019 (0.036)	0.021 (0.023)	0.022 (0.023)
Observations	5,078	144,398	144,398	144,398	144,398	184,893
R^2	0.99	0.90	0.90	0.91	1.00	1.00
Fixed effects:						
- Imp./Exp.-pair	Yes	Yes	Yes	Yes	Yes	Yes
- Imp./Exp.-time	No	No	No	Yes	Yes	Yes
- Time	Yes	Yes	Yes	No	No	No

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$
Note: Standard errors clustered at the imp./exp.-pair level.

Table A1: Regressions (non-balanced trade), 1992-2002

Dep. variable: Log imports (PPML: imports)	(1) Sml. Sample, 1992-2013	(2) World Sample, 1992-2013	(3) World Sample, 1992-2013	(4) World Sample, 1992-2013	(5) World Sample, 1992-2013 (PPML: no 0s)	(6) World Sample, 1992-2013 (PPML)
Euro	0.102 (0.031)***	-0.050 (0.033)	-0.095 (0.034)***	-0.063 (0.051)	-0.018 (0.036)	-0.013 (0.036)
Log product of GDPs	0.876 (0.310)***	0.762 (0.048)***	0.780 (0.048)***			
Log product of per-capita GDPs	-0.414 (0.330)	-0.135 (0.048)***	-0.152 (0.048)***			
EU	-0.050 (0.038)	0.268 (0.036)***	0.349 (0.038)***	0.411 (0.059)***	0.209 (0.055)***	0.202 (0.055)***
EU trend	0.010 (0.003)***	0.009 (0.002)***	0.011 (0.002)***	0.042 (0.004)***	0.010 (0.003)***	0.011 (0.003)***
FTA			0.149 (0.023)***	0.180 (0.026)***	0.064 (0.026)**	0.063 (0.026)**
Other com. curr.			-0.135 (0.028)***	-0.019 (0.038)	0.022 (0.029)	0.023 (0.029)
Observations	10,160	314,620	314,620	314,620	314,620	390,393
R^2	0.98	0.88	0.88	0.89	0.99	0.99
Fixed effects:						
- Imp./Exp.-pair	Yes	Yes	Yes	Yes	Yes	Yes
- Imp./Exp.-time	No	No	No	Yes	Yes	Yes
- Time	Yes	Yes	Yes	No	No	No

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$
Note: Standard errors clustered at the country-pair level.

Table A2: Regressions (non-balanced trade), 1992-2013