

A union divided? The euro and trade in the core and the periphery

Joseph Kopecky*

15th November 2025

Abstract

Has the euro improved trade evenly across member states? This paper revisits the impact of the euro on trade, focusing on systematic heterogeneity between core and peripheral members. I develop a stylized model linking trade flows to two mechanisms implied by monetary integration: the elimination of bilateral exchange-rate volatility and potential relative price-level convergence upon entry. The model predicts heterogeneous effects depending on structural characteristics of member economies. Using bilateral trade data from 1960–2018, I estimate a gravity model with Poisson pseudo maximum likelihood and apply a doubly robust inverse-propensity score weighting estimator. I find that the average EMU effect masks substantial heterogeneity across member states. On average, euro membership increases trade by about 6%, but trade originating from core members increased by an estimated 12%. Trade between periphery members is estimated to increase by a similar amount, but flows from the periphery to core decline by an estimated 7%, consistent with predictions from the theoretical model.

*Trinity College, Dublin (jkopecky@tcd.ie).

1. INTRODUCTION

It was the best of times, it was the worst of times...

- Charles Dickens, *A Tale of Two Cities*

The twenty-first century began with eleven European countries binding their economic fortunes in one of the boldest macroeconomic experiments in modern history. The first decade of the European Monetary Union (EMU) was a spring of hope, featuring: low inflation, relative business cycle stability, and further integration among an expanding group of member states. The period since the global financial crisis might be better characterized as a winter of despair, with financial instability, asymmetric sovereign risk, and the return of inflation. This period has created a unique set of challenges for policymakers and highlighted the critical differences between core and periphery members of the union. A key benefit of eurozone membership is the role that the common currency plays in promoting trade between members. Early literature provided extreme optimism with large, if implausible, estimates of the size of these benefits. Recent work suggests a positive, but modest (with unclear statistical significance) effect of the EMU on these trade flows. In this paper, I explore the extent to which the euro affects trade differently across members belonging to its so-called core and periphery. I find that the data tell a tale of two eurozones: with core countries enjoying large improvements to their export flows to both their fellow core member states as well as periphery members of the bloc, while those on the periphery experience improvements in their trade with each other, but a decrease in exports to the core. This finding can be rationalized in an Armington-style model if the euro has not only eliminated exchange rate volatility, which should improve trade across the board, but has also led to an appreciation of periphery prices relative to their core trading partners.

The literature estimating the effect of currency unions on trade is large and contentious. It stems from the influential work of [Rose \(2000\)](#), who found a doubling of trade among countries sharing a common currency. Early estimates studying the eurozone suggest a much smaller, if still eye-popping, effect with [Glick and Rose \(2016\)](#) finding a 54% increase. Much work has been done in this area that has tempered those expectations, with a recent meta-analysis by [Polák \(2019\)](#) suggesting that the EMU trade effect is likely somewhere between 2–6%. [Rose \(2017\)](#) when considering the large distribution of euro-effects among empirical papers, makes a case for those finding larger estimates. He points out that studies with small (or zero) estimates of the EMU on trade rely predominantly on truncated samples, usually restricting analysis to a subset of high income countries. He argues that such truncation biases estimates downward, though the mechanism is not clear. [Kopecky](#)

(2024) shows that such selection of the estimation sample actually provides an ad-hoc way of *improving* the comparability of “treated” EMU countries with the “control” group. When measuring the differences of means between these groups on observable characteristics, such ad hoc samples are more comparable and thus should provide more reliable estimates. His work suggests propensity score approaches that deliver such comparability in a data-driven way, through a first stage selection into monetary union treatment as a way to both truncate the sample and weight second stage estimates. In the preferred specification using such an estimator, [Kopecky \(2024\)](#) suggests a modest, but significant 5-7% boost in trade attributed to EMU membership. Here I use this estimate to replicate a similarly sized (and marginally significant) effect, with the euro increasing trade by 6%.

The contribution of this paper is to explore the potentially heterogeneous effects of eurozone membership across core and periphery members. While euro members have undergone a process of convergence in order to join the union, they differ in tangible ways. Leveraging propensity score methods, I show that much of the trade benefit experienced by periphery countries disappears when improving the comparability of treated and control groups. There is ample evidence that aggregation of individual unions into a single currency union estimate masks substantial heterogeneity. Earlier work such as [Campbell \(2013\)](#) and [Campbell and Chentsov \(2023\)](#) demonstrates the importance of historical context and existing dynamic trends driving much of the large estimates found in the literature on currency unions. In the context of a broad range of currency unions, [Kopecky \(2023\)](#) uses the same propensity score estimates to show that individual currency union effects vary substantially when looking at disaggregated currency unions in isolation. I take this idea one step further, suggesting that aggregation to the EMU of various member experiences masks substantial underlying differences. To my knowledge, this is the only paper to do so in the context of the eurozone.

The “core–periphery” lens in the EMU typically refers to persistent structural differences in: economic size, productivity, and institutions. This has often been studied in the context of asymmetric cyclical behavior across member states to understand challenges related to monetary policy implementation. One such paper is [Bayoumi and Eichengreen \(1992a\)](#) who identified clear core–periphery clusters in Europe as a challenge for the implementation of joint monetary policy. Subsequent research by [Arestis and Sawyer \(2011\)](#) shows that structural differences remain important as a challenge for joint European policy. While I find asymmetric effects of the EMU on trade between core-and periphery I suggest that such an effect is consistent with EMU improving core-periphery integration. This is consistent with findings of [Campos and Macchiarelli \(2016\)](#) who revisit the influential work of [Bayoumi and Eichengreen \(1992b\)](#), finding that the EMU has weakened these core-periphery differences.

Most of the recent work investigating core-periphery divides focus on *business-cycle synchronization* and *financial integration* with a large literature finding that business cycles co-move more tightly within the core than between core and periphery members, with these asymmetries especially pronounced around the global financial crisis and the sovereign debt crisis (Belke, Domnick, and Gros, 2017; Mink, Jacobs, and de Haan, 2007; Montoya and de Haan, 2008). On the financial side, persistent current account imbalances and asymmetric financing patterns between core and periphery members have been well documented (De Santis and Cesaroni, 2016; Lane and Pels, 2012; Atoyan, Manning, and Rahman, 2013). Related work on the euro’s “winners and losers” highlights heterogeneous welfare consequences across member states (Puzzello and Gomis-Porqueras, 2018; Fernández-Villaverde, Garicano, and Santos, 2013), reinforcing the view that integration has been uneven across multiple dimensions. By contrast, the literature estimating the euro’s impact on *trade* has largely focused on average effects.

A few contributions emphasize heterogeneity indirectly, for example, by examining early versus late EMU entrants, trade creation versus diversion patterns, or differences by currency union cohorts. Work such as Baldwin, Skudelny, and Taglioni (2005) and Flam and Nordström (2007) drill down into sectoral level flows finding potential differences in euro effects at that level, with Flam and Nordström (2007) pointing to increased vertical specialization as a result of currency union membership. Badinger and Breuss (2006) explore the role of country size on the relative gain of currency union members, suggesting that large countries benefit more, but then in their subsequent test of this effect within the eurozone in Badinger and Breuss (2009) find that small countries benefit more from euro membership. However, explicit disaggregation between *core* and *periphery* members in this context remains scarce, despite suggestive evidence that aggregation may mask substantial variation in bilateral trade responses.

This paper brings the core–periphery perspective squarely to the trade channel by estimating separate euro effects for core-core, core-periphery, and periphery-periphery country pairs. I first provide a conceptual framework for heterogeneous euro effects in section 2. I then describe my empirical model and doubly robust gravity equation estimator in section 3. Empirical estimates are presented in section 4 and section 5 concludes.

2. CONCEPTUAL FRAMEWORK

To ground empirical estimates I begin by developing a stylized Armington-style trade model with two types of countries: *core* (C) and *periphery* (P). This helps formalize two mechanisms through which currency union (CU) membership can affect trade flows differentially across

pairs: (i) reductions in bilateral exchange-rate volatility, and (ii) other fundamental factors that shift relative prices.

2.1. Setup

Consider a world with two groups of countries, denoted C (core) and P (periphery). Each country i produces a differentiated variety that can be traded internationally subject to iceberg trade costs $\tau_{ij} \geq 1$. Consumers in destination j have CES preferences over varieties indexed by their origin i :

$$X_{ij} = E_j \left(\frac{p_{ij}}{P_j} \right)^{1-\sigma}, \quad P_j = \left(\sum_m p_{mj}^{1-\sigma} \right)^{\frac{1}{1-\sigma}}, \quad (1)$$

where E_j denotes total expenditure in j , p_{ij} is the consumer price of the variety from i , P_j the corresponding CES price index, and $\sigma > 1$ the elasticity of substitution across varieties.

Under monopolistic competition and CES demand, firms charge a constant markup over marginal cost:

$$\mu = \frac{\sigma}{\sigma - 1}. \quad (2)$$

Let a_i denote the unit input requirement in country i (the inverse of productivity), w_i the nominal wage, and e_{ij} the nominal exchange rate—defined as the price of currency i in terms of currency j . The marginal cost of supplying market j in j 's currency is therefore $a_i w_i \tau_{ij} e_{ij}$. The optimal export price is then

$$p_{ij}^j = \underbrace{\mu a_i \tau_{ij}}_{\text{technology / markup}} \times \underbrace{w_i}_{\text{producer costs}} \times \underbrace{e_{ij}}_{\text{nominal exchange rate}} \times \underbrace{\kappa(\sigma_{ij}^2)}_{\text{exchange rate volatility adjustment}}, \quad (3)$$

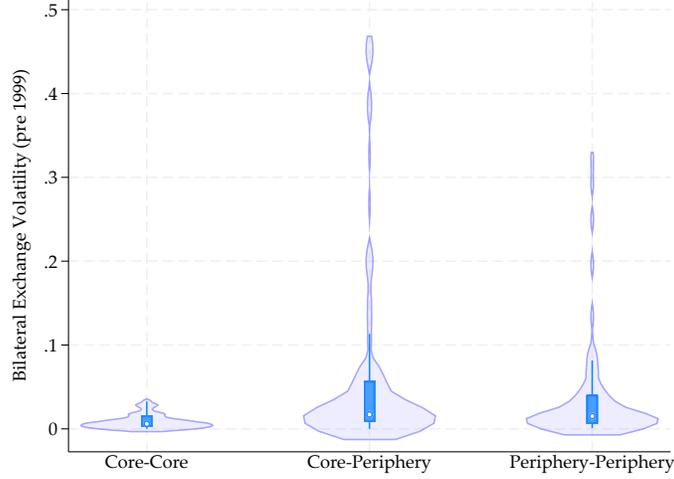
where $\kappa'(\sigma_{ij}^2) > 0$ and $\kappa(0) = 1$. The function $\kappa(\sigma_{ij}^2)$ captures the convexity adjustment that arises when firms set prices before exchange rates are realized: with log-normal exchange-rate shocks, $\kappa(\sigma_{ij}^2) = \exp(\frac{1}{2}\sigma_{ij}^2)$, so higher volatility raises preset prices in foreign currency. When bilateral volatility collapses, as under a currency union, $\kappa(\sigma_{ij}^2) \rightarrow 1$, reducing the foreign-currency price even if local costs remain unchanged. Under log-normal exchange-rate shocks, the reduction in volatility with EMU implies $\Delta \ln \kappa(\sigma_{ij}^2) = -\frac{1}{2}\sigma_{ij}^{2,\text{pre}}$.

Taking logs and total differentials of (3) yields

$$\Delta \ln p_{ij}^j = \Delta \ln a_i + \Delta \ln w_i + \Delta \ln \tau_{ij} + \Delta \ln e_{ij} + \Delta \ln \kappa(\sigma_{ij}^2), \quad (4)$$

showing that any change in technology, producer costs, trade frictions, nominal exchange

Figure 1: *Relative Exchange Rate by Core-Periphery Pairs*



Periphery-core (PC) is identical to core-periphery and thus omitted here.

rates, or exchange-rate volatility can affect relative prices.

2.2. Currency Union Effects

We consider two mechanisms through which CU membership affects trade:

(1) Exchange-rate volatility elimination. CU membership sets $e_{ij} = 1$ and $\sigma_{ij}^2 = 0$ for intra-union pairs. The fall in volatility reduces $\kappa(\sigma_{ij}^2)$, lowering effective trade costs. The log change in trade flows is

$$\Delta \ln X_{ij}^{(\text{vol})} = (1 - \sigma) \Delta \ln \kappa(\sigma_{ij}^2) = (1 - \sigma) \left(-\frac{1}{2} \sigma_{ij}^{2, \text{pre}} \right) = \frac{1}{2} (\sigma - 1) \sigma_{ij}^{2, \text{pre}} > 0,$$

with the expected gain proportional to pre-currency union volatility. While the sign of this effect is positive for all bilateral trading partners, periphery-related pairs had higher volatility broadly speaking, creating a channel through which heterogeneity of a euro affect can arise. To motivate this I plot relative bilateral exchange volatility between core-core (CC), core-periphery (CP), and periphery-periphery (PP) partners from 1980 through 1998 in Figure 1. This raw volatility simplifies an admittedly complex transition among many of these partners, through various floats and fixed exchange regimes, but motivates that we might expect differences across each type of core and periphery partnership.

The finding in Figure 1 that CP/PC volatility is higher than PP volatility is somewhat puzzling as the stylized model, if one assumes higher periphery volatility than core pre-

EMU would predict a strict ordering of:

$$\sigma_{CC}^{2,\text{pre}} < \sigma_{CP}^{2,\text{pre}} = \sigma_{PC}^{2,\text{pre}} < \sigma_{PP}^{2,\text{pre}}. \quad (5)$$

However, the difference in these groups is relatively small and could possibly be explained by stronger co-movement among periphery members.

(2) Relative price-level shifts. Joining the currency union may induce a real appreciation of one trading partner relative to another. In this case we assume there is an appreciation of periphery prices relative to those in the core. Letting $\Delta \ln w_P > \Delta \ln w_C$, representing a larger price-level increase in the periphery. This raises p_{PC} relative to p_{CP} , implying

$$\Delta \ln X_{PC} = (1 - \sigma)\Delta \ln p_{PC} < 0, \quad \Delta \ln X_{CP} = (1 - \sigma)\Delta \ln p_{CP} > 0.$$

There is no first-order effect on X_{CC} and X_{PP} from relative prices, as both sides in each pair type move symmetrically. This of course can come through any of the remaining components of Equation 4 (i.e.: w_i, a_i, τ_{ij}). To again motivate that there is a strong degree of price integration throughout this period, I plot the relative export prices for these core and periphery groups in Figure 2. These data come from version 11 of the Penn World Tables (Feenstra, Inklaar, and Timmer, 2015). The top two panels show the relative export prices of each core country averaged over their core and periphery counterparts. While the bottom two show the same for periphery. Bold lines represent group averages among the group which are by definition zero among CC and PP pairs. We see an appreciation of periphery prices relative to core evident in both the Figure 2b and Figure 2c. Additionally all four show convergence of prices toward trends. While some of these prices and volatilities may come from pure exchange rate effects, or be driven by other sources of European integration, the continued appreciation and convergence post-EMU entry suggests potential price convergence through one of these other channels.

To summarize the implications of this through both channels I consider that some price increase is driven by $\Delta \ln w_P > \Delta \ln w_C$, though in principle this could be any of the other non-exchange rate components of Equation 4. Then we can say the following.

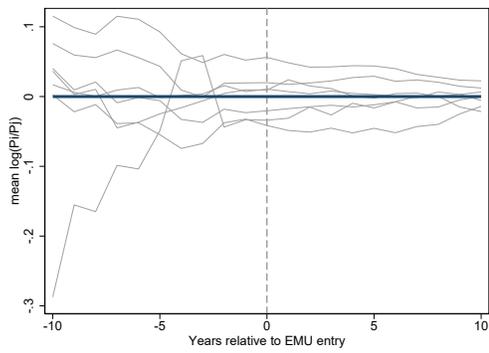
Proposition 2.1. *Suppose that CU membership leads to (i) a reduction in bilateral exchange-rate volatility, $\sigma_{ij}^2 \rightarrow 0$, with $\sigma_{ij}^{2,\text{pre}}$ larger for periphery-related pairs than for core-core pairs and (ii) a relative price increase for periphery goods driven by, $\Delta \ln w_P > \Delta \ln w_C$. Let $\sigma > 1$. Then:*

1. *The log change in bilateral trade flows for pair (i, j) is*

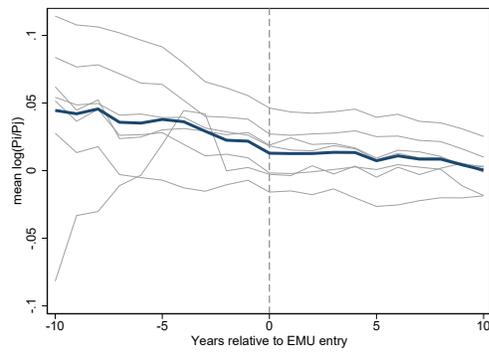
$$\Delta \ln X_{ij} = (1 - \sigma) \Delta \ln p_{ij} = (1 - \sigma) [\Delta \ln w_i + \Delta \ln \kappa(\sigma_{ij}^2)],$$

Figure 2: Aggregate Relative Export Prices

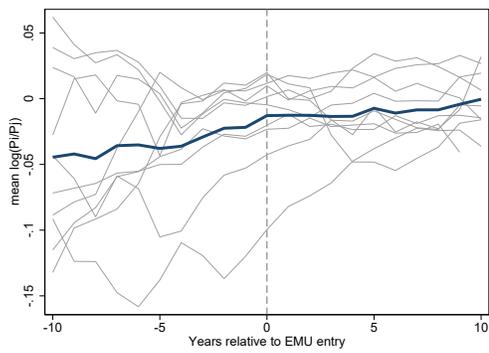
(a) Core: Average over Core Partners



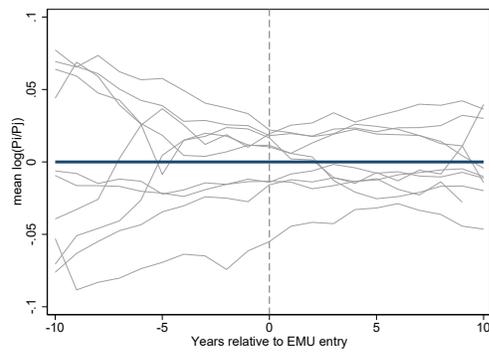
(b) Core: Average over Periphery Partners



(c) Periphery: Average over Core Partners



(d) Periphery: Average over Periphery Partners



where $\Delta \ln \kappa(\sigma_{ij}^2) = -\frac{1}{2}\sigma_{ij}^{2,\text{pre}}$.

2. The sign of $\Delta \ln X_{ij}$ by pair type is given in [Table 1](#).

Table 1: Predicted Sign of Trade Effects by Pair Type

Pair type	Relative price effect	Volatility effect	Overall trade effect
Core \rightarrow Core (CC)	0	+	+
Core \rightarrow Periphery (CP)	+	+	+
Periphery \rightarrow Core (PC)	-	+	Ambiguous
Periphery \rightarrow Periphery (PP)	0	+	+

This simple structure captures the asymmetric effects of CU membership on different types of bilateral flows as summarized by this theoretical model. Further, we can generate some expectations about the relative size of the effect both from [Table 1](#) and magnitudes in [Figure 1](#). Core-core flows experience modest gains given the model mechanism (which assumes symmetry of all core countries) and have no relative price channel outside of the impact of reduced volatility. One might expect gains to be modest given the relatively small pre-1999 bilateral exchange rate volatility for these members in [Figure 1](#). Core-periphery flows likely benefit from both volatility reduction and may expect that core export prices become more favorable in periphery due to convergence of other non-exchange rate fundamentals. Symmetrically, exports from periphery-to-core face offsetting forces. Exports between periphery pairs see positive impacts only through the exchange volatility channel. Of course there may be within group differences that are important, with [Figure 2](#) suggesting that periphery countries perhaps have much wider price dispersion, which does appear to be converging over time, potentially opening channels not discussed here through which PP pairs might improve their trade with one another. Indeed I will find similar, and positive effects among C-C, C-P, and P-P flows, though statistical significance is somewhat in line with these expected magnitudes.¹ I find that P-C flows appear to be weakly negative in line with what one might expect from this framework.

3. DATA AND METHODOLOGY

Bilateral export data come from the CEPII gravity dataset, ([Conte, Cotterlaz, and Mayer, 2021](#)). I construct an updated version of the [Glick and Rose \(2016\)](#) EMU membership to include new entrants.² Non-EMU currency unions are dropped as including them could

¹As we will see core-to-core coefficient is less precisely estimated, while P-P and C-P effects are quite robust.

²I also exclude small French territories that are part of the [Glick and Rose \(2016\)](#) EMU variable.

bias results. Similarly, estimations on subsets of EMU members exclude all other EMU trading partners.³

I classify euro members into *core* and *periphery* following widely used groupings in the EMU literature (e.g., [Bayoumi and Eichengreen, 1992a](#); [Campos and Macchiarelli, 2016](#)). The complete list used in this paper is reported in [Table 2](#).

Table 2: *Baseline Core–Periphery Classification (Euro Members)*

Core	Periphery
Austria; Belgium; Finland; France; Germany; Luxembourg; Netherlands	Cyprus; Greece; Italy; Ireland; Latvia; Lithuania; Malta; Portugal; Slovakia; Slovenia; Spain

I specify the gravity equation for trade including rich set of exporter-year and importer-year fixed effects to account for structural multilateral resistance terms, as suggested in [Head and Mayer \(2014\)](#). In addition, I include time-invariant dyadic fixed effects that are shown to have been important in the context of currency union estimates by [Baldwin and Taglioni \(2007\)](#). This is given by:

$$X_{ijt} = e^{\gamma D_{ijt} + \beta Z_{ijt} + \phi_{ij} + \psi_{jt} + \lambda_{it}} \quad (6)$$

where X_{ijt} are exports from country i to country j at time t . The policy treatment of interest is, EMU_{ijt} , representing euro membership for the country-pair in year t . I will show average euro membership, but also separately estimate euro membership for core and periphery members, first to all other members of the monetary union and then to other core or periphery destinations. Any additional time-varying controls are captured by Z_{ijt} while exporter-time, importer-time, and dyadic fixed effects are given by: λ_{it} , ψ_{jt} , and ϕ_{ij} respectively.

The log-linearization of [Equation 6](#) is vulnerable to bias in the presence of heteroskedasticity, a problem shown to be quantitatively important in the context of trade by [Silva and Tenreyro \(2006\)](#) and [Silva and Tenreyro \(2011\)](#). As a solution, they suggest a Poisson pseudo-maximum likelihood (PPML) estimator of [Equation 6](#) directly, which is robust to this particular source of bias. In addition to correcting this, PPML estimation also allows for inclusion of the large set of zero-flows present in bilateral trade data. Generally I find little effect if missing data is treated as a zero flow or excluded so in specifications with the full data for PPML I opt to include missing trade flows as zero.⁴ Generally log specifications

³In both cases this is because there may be some currency union (or EMU) effect common among groups even if there are also asymmetric effects.

⁴As I will discuss at length my preferred estimates drop the majority of observations in any case so this is

are qualitatively similar in my preferred estimates, suggesting that at least some of this bias is solved with my propensity score truncation and weighting.

3.1. Inverse Propensity Score Weighting: A Doubly Robust Estimator

In the early literature estimating empirical currency union effects, [Persson \(2001\)](#) showed that matching estimators substantially reduce implied trade effects found by [Rose \(2000\)](#). These estimates, along with similar EMU estimates in [Chintrakarn \(2008\)](#), use pure matching estimators⁵ that rely on correct specification of a first stage selection equation into currency union treatment. In short, such estimators attempt to use propensity scores to generate a treatment and control sample with matching characteristics and thus transform non-experimental data into something that looks as if its treatment was randomly assigned. Responses to [Persson \(2001\)](#) by [Rose \(2001\)](#) rightly point out that the model of first-stage selection performs relatively poorly in terms of fit and out-of-sample predictive power, making it difficult to judge their merit relative to the well-studied gravity approach. The critiques put forward in [Rose \(2001\)](#) are correct in pointing out the flaws of matching estimators, but do not address the problems of selection that motivated the work of [Persson \(2001\)](#) and [Kenen \(2002\)](#).

While propensity score based estimators are not a silver bullet to identification they can be useful in terms of ensuring balance between treated and control groups and offer insight into how to improve traditional gravity estimators in a data-driven way. A class of *doubly robust* estimators use propensity scores to jointly model the selection into treatment and the outcome model. I follow [Kopecky \(2024\)](#) by using inverse propensity score weighting with regression adjustment (IPWRA). To do so I must first estimate a model of selection into currency union to generate probabilities of treatment. These are then used to both select an appropriate comparison group (ie truncate the sample), and then reweight second stage estimates of the gravity equation in [Equation 6](#). In the context of the initial debate around currency unions this is using the probability estimates of [Persson \(2001\)](#), while taking advantage of the well studied, and theoretically motivated gravity equation. The *doubly-robust* nature of the IPWRA estimator, is discussed in: [Imbens \(2004\)](#), [Lunceford and Davidian \(2004\)](#), and [Wooldridge \(2007\)](#) but the simple takeaway is that estimates are consistent if *either* model is correctly specified. As in [Kopecky \(2024\)](#) I follow notation of [Jordà and Taylor \(2016\)](#), who to my knowledge first brought this estimator to a macroeconomic context. The IPWRA estimator is given by the following equation.

only relevant for the full sample comparison, and not in my preferred specification.

⁵In particular *nearest-neighbor* and stratification approaches to matching

$$\widehat{ATE}_{IPWRA} = \frac{1}{n_1^*} \sum \left[\frac{EMU_{ij,t} (m_1 (Z_{ij,t}, \hat{\gamma}))}{\hat{p}_{ij,t}} \right] - \frac{1}{n_0^*} \sum \left[\frac{(1 - EMU_{ij,t}) (m_0 (Z_{ij,t}, \hat{\gamma}))}{1 - \hat{p}_{ij,t}} \right] \quad (7)$$

Where $EMU_{ij,t}$ is a dummy representing whether a country-pair are in the EMU at time t . This will include a baseline estimate of the EMU effect similar to [Kopecky \(2024\)](#), as well as estimates that include core and periphery flows independently. The first-stage probability of treatment is given by $\hat{p}_{ij,t}$. This estimator allows for any general second stage model for the conditional mean associated with being in the treatment group. This is denoted as: $m_{0/1} (Z_{ij,t}, \hat{\gamma})$.⁶ In this context the mean is estimated from [Equation 6](#), where $\hat{\gamma}$ is the vector of estimated regression coefficients coming from the PPML estimation of [Equation 6](#). In principle one may estimate $\hat{\gamma}$ separately across treatment and control groups, as described [Śłoczyński and Wooldridge \(2018\)](#). In the context of the gravity equation for the EMU this is not possible due to the high dimension of fixed effects suggested in [Head and Mayer \(2014\)](#), as there are more of these effects than the number of treated observations. As a result I am limited to imposing the assumption that the coefficients for estimation of these conditional means are identical across these groups.⁷ Following work of [Hirano and Imbens \(2001\)](#) and [Imbens \(2004\)](#) I normalize weights so that $n_1^* = \sum \frac{EMU}{\hat{p}}$ and $n_0^* = \sum \frac{1-EMU}{1-\hat{p}}$, to ensure that probability weights sum to one.

3.2. Modeling Selection Into Currency Unions

Estimation of [Equation 7](#) requires a first stage model of selection into EMU treatment. I specify a logistic model with for the probability that a country i is in a monetary union with a trading partner j in time t given by: $EMU_{ij,t}$. Doing this for the entire EMU models this simply as euro area membership, but to estimate the impact separately for core and periphery I consider seven separate specifications where this variable is defined as Core/Periphery exporter relative to three separate destination partners: all other EMU members, core, and periphery. In principle, one could use the same first stage estimation for the full EMU to estimate [Equation 7](#) for each of these potential trading partners in the second stage. I will show below why this is likely a poor choice. In addition to a number of standard gravity controls I include differences between origin and destination GDP, GDP per capita, and growth rates. I do this to capture not only the standard gravity effects for

⁶For ease of notation I refer to the control set as simply $Z_{ij,t}$, but note that this includes in it the rich set of fixed effects in [Equation 6](#).

⁷This is of course standard for gravity estimates, I simply note that it is not a required assumption if one were to prefer a different estimate of the conditional mean in the second stage.

economic variables, but also to capture the *relative* size between origin and destination. This may be important if there are differences between currency union arrangements between similarly sized countries and those where there are large asymmetries. My first-stage model is shown by [Equation 8](#).

$$\begin{aligned}
 EMU_{ij,t} = & \theta_0 + \theta_1 \ln Y_{i,t} \times Y_{j,t} + \theta_2 \ln y_{i,t} \times y_{j,t} + \theta_3 \ln Dist_{ij,t} \\
 & + \theta_4 \ln |Y_{i,t} - Y_{j,t}| + \theta_5 \ln |y_{i,t} - y_{j,t}| + \theta_6 |g_{i,t} - g_{j,t}| \\
 & + X_{i,j,t}
 \end{aligned} \tag{8}$$

The first three terms, are standard gravity equation variables: output ($Y_{i/j,t}$), income ($y_{i/j,t}$), and distances⁸ ($Dist_{ij,t}$) between the two countries. As discussed, I also include differences in output, per-capita GDP, and GDP growth rates ($g_{i/j}$) to capture the importance of the relative size of exporters and importers. I also include a number of available controls (X): population in origin and destination, distances, common legal traditions (both pre and post transition), and common colonial dependency.

4. RESULTS

4.1. First stage selection into EMU

Before moving to estimates of the EMU effect on trade my approach requires first stage estimation of probability of entering into the EMU. The aim is to generate a first stage model that suitably accounts for the differences in observable factors between not just eurozone and non-eurozone country pairs, but also between core and periphery pairs. To do this I specify seven different treatment groups: all EMU, core-to-EMU, periphery-to-EMU, core-to-core, core-to-periphery, periphery-to-periphery, and periphery-to-core. To understand how the monetary union affects trade between periphery members and across core and periphery members I wish to have a first stage that can flexibly allow that such bilateral pairs may be substantially different. For example many periphery EMU countries are much smaller and poorer than core countries, so that the appropriate comparison group for Germany and Slovenia, maybe not be the same as that which is best for Germany and France, or for Slovenia and Greece. In [Table 3](#) I report the results for the full first stage for each of these groups.

While the point estimates in [Table 3](#) are not the focus of this paper there are a few results that are telling with respect to heterogeneity across core and periphery pairs. First, the model fit (pseudo R^2) is highest for core-to-core pairs, and lowest for periphery-to-

⁸I use the population weighted distance between the two most populous cities in each country.

Table 3: First-Stage Logit Estimations

	(1) EMU	(2) C-EMU	(3) P-EMU	(4) C-C	(5) C-P	(6) P-P	(7) P-C
$\ln(GDP_1 - GDP_2)$	0.20*** (0.02)	0.33*** (0.03)	0.18*** (0.02)	0.53*** (0.05)	0.29*** (0.03)	0.17*** (0.03)	0.29*** (0.03)
$\log(GDPpc_1 - GDPpc_2)$	-0.23*** (0.01)	-0.17*** (0.02)	-0.24*** (0.02)	-0.72*** (0.03)	0.22*** (0.03)	-0.43*** (0.02)	0.22*** (0.03)
$\log(GDP_1 \times GDP_2)$	0.11*** (0.01)	0.15*** (0.02)	0.10*** (0.01)	0.04 (0.03)	0.17*** (0.02)	0.08*** (0.02)	0.17*** (0.02)
$\log(GDPpc_1 \times GDPpc_2)$	1.13*** (0.02)	1.23*** (0.03)	0.97*** (0.02)	1.87*** (0.07)	1.00*** (0.03)	0.89*** (0.03)	1.00*** (0.03)
$ Growth_1 - Growth_2 $	-0.16*** (0.01)	-0.23*** (0.01)	-0.13*** (0.01)	-0.38*** (0.02)	-0.18*** (0.01)	-0.10*** (0.01)	-0.18*** (0.01)
Pop. Origin	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.00*** (0.00)	-0.01*** (0.00)	-0.02*** (0.00)
Pop. Destination	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.02*** (0.00)	-0.01*** (0.00)	-0.00*** (0.00)
Distance	-5.57*** (0.11)	-6.33*** (0.20)	-4.93*** (0.12)	-12.73*** (0.68)	-5.29*** (0.20)	-4.48*** (0.15)	-5.29*** (0.20)
Common leg. (pre tran.)	0.39*** (0.07)	0.35** (0.11)	0.43*** (0.09)	1.25*** (0.33)	0.41*** (0.12)	0.60*** (0.11)	0.41*** (0.12)
Common Leg.	0.04 (0.07)	0.07 (0.11)	0.03 (0.08)	-1.38*** (0.32)	0.26* (0.11)	-0.30** (0.11)	0.26* (0.11)
Col. dep.	-0.92*** (0.13)	-0.42** (0.16)	-1.42*** (0.22)	0.45 (0.24)	-0.79*** (0.23)		-0.79*** (0.23)
Observations	1524294	1521683	1522215	1520486	1520801	1505516	1520801
Pseudo R^2	0.629	0.656	0.535	0.749	0.558	0.464	0.558
p_{min}	0.005	0.005	0.005	0.007	0.005	0.005	0.005
p_{max}	0.875	0.853	0.667	0.937	0.533	0.444	0.534

All estimates use logistic regression on a binary outcome reflecting eurozone "treatment". EMU uses all bilateral EMU pairs as the outcome variable. The remaining titles refer to Origin-Destination classification where C refers to core EMU members and P refers to periphery.

periphery pairs. This is reflective of the fact that core EMU countries have much more in common on observable characteristics than periphery countries and econometric model is able to do a better job identifying them from observables, while struggling more with periphery-to-periphery pairs. One can further see some of the, expected, asymmetries when comparing estimated coefficients. EMU countries are much larger than the average bilateral pair in the sample in terms of their GDP and GDP per capita, though this is much stronger for core countries substantially than periphery. Pairs are generally farther apart in size (the absolute difference of GDP) than the average pair bilateral trading pair in the sample. EMU countries are on average closer together in income (difference in per-capita GDP), but this effect reverses for core-periphery pairs, suggesting that the gap in income between the core and periphery of the EMU is larger than the average gap in bilateral trading partners in the sample.

Ultimately the goal of the models in [Table 3](#) are to improve the comparability of the observations that serve as a control set to the various EMU outcomes. To achieve this one wants to have a model that provides some information to usefully sort groups into these groups, while also having overlap between “treated” and “control” observations. I estimate all specifications only on data for which the control set falls within the support of the treatment group, reporting these maximum and minimum probabilities of treatment in [Table 3](#). To illustrate this overlap, and further motivate the need to conduct separate analysis across each of these sub-components of the EMU I plot the densities showing probability of treatment for the Column 1 estimation of [Table 3](#). These are shown in [Figure 3](#).

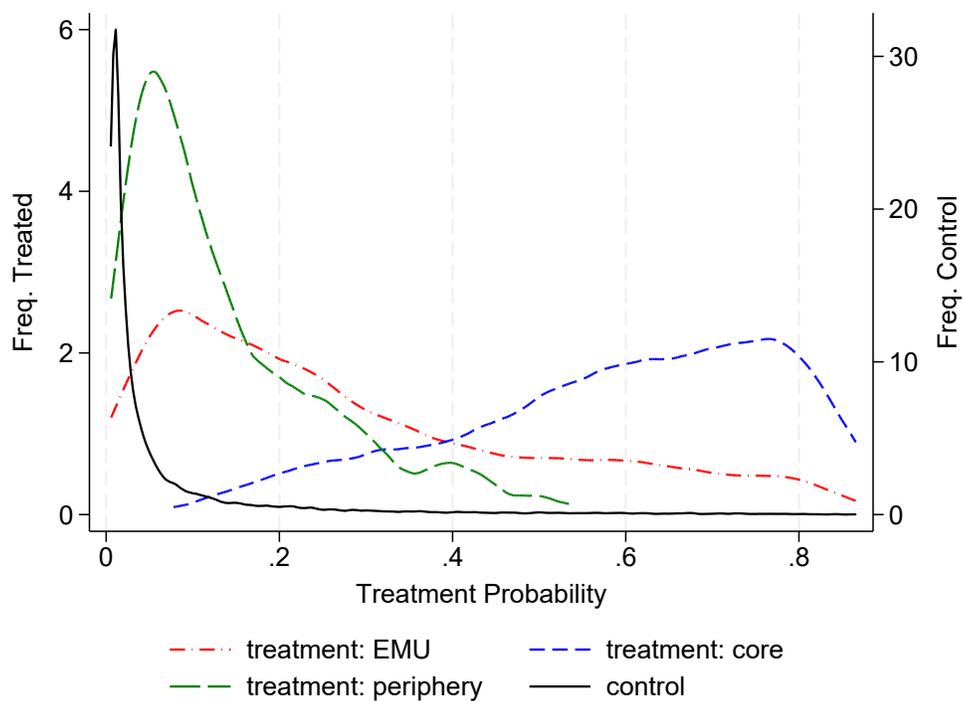
The black solid line is the control group⁹ The red dash-dot line is the EMU “treatment” that was the outcome of this particular logit model. There is considerable overlap between the two groups and the model does provide some predictive power here. However, I also show that there is a great deal of heterogeneity underlying this group by also plotting the core-core and periphery-periphery probabilities within this model.¹⁰ Here we can see that the EMU selection model provides very different looking estimates across these pairs. For this reason I run separate first stage models for each sub-EMU group.

The next relevant question is whether using the probabilities from [Table 3](#) actually accomplish the goal of creating comparable “treatment” and “control” groups for use in second stage gravity estimations. To show that this is the case, I plot the difference in means between treated and control units for all of the specifications I will later report for a

⁹Because the control group is extremely large relative to the number of EMU pairs control density is plotted on the right-hand axis to improve readability.

¹⁰For exposition I show the probabilities for these groups for the same estimation in Column 1 estimated for the full EMU, not the estimations that treat these sub-groups as the outcome themselves. These figures are interesting in their own right, and important diagnostically, but do not add a great deal and have been left out of the paper for brevity.

Figure 3: Kernel density: probability of treatment (full EMU model)



This figure plots predicted probability of treatment for all EMU countries, these are reported in Column 1 of Table 3. The logistic equation of Equation 8 is estimated on the full EMU sample (red dot-dashed line), but substantial differences exist between core (dashed blue) and periphery (dashed green) probabilities from such a model.

Figure 4: Comparability of treatment and control groups

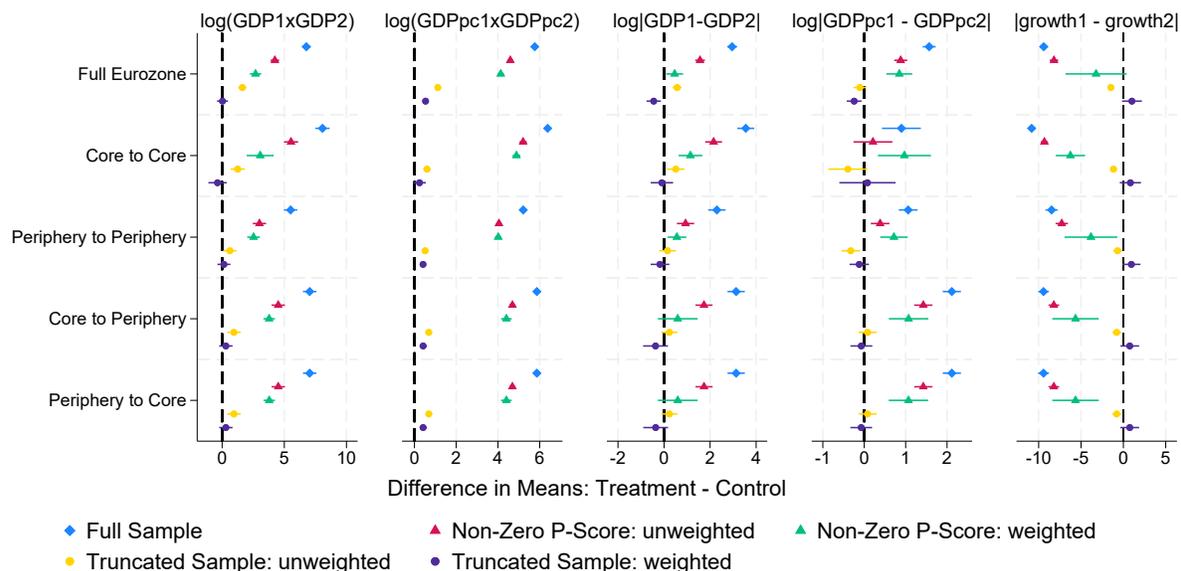


Figure compares means between treated and control groups across five variables used in the first stage [Table 3](#) estimations and five versions of “treatment” variables (the full EMU, and disaggregated EMU relationships between core and periphery countries). Bars represent 95% clustered confidence intervals.

number of the variables in [Table 3](#). These are shown in [Figure 4](#).

The initial difference in means between these EMU groups and the control sets are shown with the blue diamonds and highlight the problem that this method seeks to address. These are dramatically different. From there I show how using probabilities to both truncate the estimation sample, and weight the groups improves the balance across these observable. First I show the sample for which the first stage estimation is able to produce a non-missing probability of treatment. This is essentially the sample for which data is non-missing are satisfied. This sample (red triangles) already moves many of these much closer to zero. Weighting on this same sub-sample (green triangles) further improves fit. Finally truncation of the sample to ensure both overlap condition (treatment and control support must overlap) and ad-hoc removal of extreme values of probabilities brings differences quite close to zero, even in the absence of weighting. Adding weights on this limited sample in most cases improves fit to be nearly zero.

4.2. The Effect of the EMU in the core and periphery

In [Table 4](#) I report results for the trade impact of the EMU across each of the combinations of trading partner types. All estimations use the PPML estimator and include pair, importer-

year, and exporter-year fixed effects, along with a control for regional trade agreements. I report two versions of standard errors for each estimate. In parentheses are the most conservative three-way clustered errors by importer-year, exporter-year, and pair; in brackets are errors clustered by importer-year and exporter-year. Because euro adoption occurred for many countries in a short time window, the identifying variation is concentrated in a few years. Consequently, pair-specific shocks are highly correlated across pairs during these years, so allowing correlation both within pairs over time and across importer-year and exporter-year shocks leaves few independent clusters and thus larger, more conservative standard errors. Given the lack of consensus in the gravity literature¹¹ on the appropriate clustering dimension, reporting both provides a clear sense of estimate robustness.¹² In discussing results, when significance under two-way clustering disappears with the addition of the pair cluster, I describe this as *weak evidence*. In the preferred specification (Column 5), most results remain strong with two-way clustering, but only those for Periphery–Periphery and Core–Periphery (and weakly Core–to–EMU) survive the stricter three-way clustering. The near-zero effect for Periphery–to–EMU/Core is itself informative and consistent with the conceptual framework in [section 2](#).

Column 1 provides the standard, full sample, evidence that retains all possible bilateral trade data. Moving from left to right moves through different samples and weighting schemes, with columns 2 & 3 estimating the unweighted and weighted coefficients on the sample for which a first stage propensity score was estimated, while 4 & 5 further truncate this sample to require common support among currency union (the EMU/sub-EMU relationship) and non-union controls as well as trimming of extreme propensity score outliers. This final column is the one with closest comparability between these *treated* and *control* groups as earlier shown in [Figure 4](#), making these my preferred estimates. The headline full EMU estimates are in line with the literature, starting at 8.5% in the full sample and declining only slightly to 6.3% when using the doubly robust IPWRA estimator on the truncated sample. This lower estimate is significant only when using the less conservative two-way clustering. Though not new, this result adds to a large body of evidence of a modest, and weakly significant trade effects. In the meta analysis by [Polák \(2019\)](#) these effects are found to be 2-6% and using the same estimator¹³ [Kopecky \(2024\)](#) reports between

¹¹Many recent papers use only importer-year and exporter-year clusters.

¹²I do not report the less conservative heteroskedastic-robust estimates, common in many EMU trade papers, though they remain strongly significant for my preferred specification except in the case of the near-zero result for P–EMU. Guidance from [Abadie, Athey, Imbens, and Wooldridge \(2023\)](#) suggests that more conservative clustered standard errors are not automatically preferable when the data represent a population rather than a random sample. I found the discussion in [McKenzie \(2017\)](#) of [Abadie et al. \(2023\)](#) particularly useful, especially its analogy to [Imbens \(2004\)](#).

¹³I use the exact same estimator, though a first stage that uses slightly fewer controls to ensure that I keep sufficient overlap in sub-EMU groups.

Table 4: Gravity Equation: PPML Estimates

	(1)	(2)	(3)	(4)	(5)
EMU All	0.085 (0.041)** [0.018]***	0.074 (0.041)* [0.018]***	0.076 (0.040)* [0.019]***	0.037 (0.056) [0.025]	0.063 (0.064) [0.029]**
r2	0.991	0.991	0.992	0.996	0.996
N	1,521,887	710,563	710,563	49,783	49,783
Core to EMU	0.111 (0.049)** [0.021]***	0.097 (0.048)** [0.021]***	0.101 (0.048)** [0.021]**	0.110 (0.074) [0.035]***	0.139 (0.084)* [0.042]***
r2	0.991	0.991	0.992	0.997	0.997
N	1,519,276	707,602	707,602	23,527	23,527
Periphery to EMU	0.063 (0.052) [0.021]***	0.058 (0.052) [0.021]***	0.053 (0.049) [0.021]**	-0.015 (0.062) [0.025]	-0.009 (0.062) [0.025]
r2	0.990	0.990	0.991	0.995	0.996
N	1,519,808	708,484	708,484	45,854	45,854
Core to Core	0.119 (0.064)* [0.028]***	0.103 (0.064) [0.027]***	0.097 (0.062) [0.028]***	0.108 (0.095) [0.052]**	0.127 (0.096) [0.054]**
r2	0.991	0.991	0.992	0.996	0.997
N	1,518,079	705,979	705,979	6,442	6,442
Periphery to Periphery	0.246 (0.060)*** [0.029]***	0.203 (0.054)*** [0.027]***	0.198 (0.049)*** [0.026]***	0.137 (0.067)** [0.027]***	0.129 (0.062)** [0.026]***
r2	0.990	0.990	0.991	0.994	0.995
N	1,518,611	692,942	692,942	35,864	35,864
Core to Periphery	0.104 (0.054)* [0.020]***	0.095 (0.053)* [0.021]***	0.104 (0.052)** [0.021]***	0.108 (0.053)** [0.026]***	0.119 (0.055)** [0.027]***
r2	0.990	0.990	0.992	0.996	0.997
N	1,518,394	707,066	707,066	23,090	23,090
Periphery to Core	-0.013 (0.058) [0.021]	-0.023 (0.057) [0.021]	-0.026 (0.054) [0.021]	-0.085 (0.073) [0.028]***	-0.074 (0.075) [0.030]**
r2	0.990	0.990	0.992	0.996	0.996
N	1,518,394	707,062	707,062	23,037	23,037
Sample: Full	✓				
Sample: p-score non-missing		✓	✓		
Sample: p-score truncated				✓	✓
Unweighted	✓	✓		✓	
IPWRA			✓		✓

Each panel reports the PPML gravity equation estimates for eurozone membership from Equation 6 for bilateral exports between member countries. All estimations include pair, exporter-year, and importer year fixed effects and a control for regional trade agreements. Two clustered standard errors are reported. In parenthesis I show the more conservative three-way clustered by: pair, exporter-year, and importer year. Brackets report two-way: importer-year and exporter-year clustered errors. Significance of these errors are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5-7% increase in trade associated with the EMU.

The next two sets of estimates (the second and third block of [Table 4](#)) show estimated EMU effects of core and periphery exporters trade to all EMU destinations. These two estimates already suggest some degree of core-periphery heterogeneity. This difference looks small in the full-data estimate, but shows core exports improving by nearly five percentage points relative to periphery. Improving the comparability of the two groups across columns slightly reinforces this positive effect for core exporters, while completely eroding any estimated trade boost on the periphery. The motivating theory above suggested an unambiguously positive impact for core exporters, but potentially ambiguous effects on the periphery where flows between periphery pairs were predicted to be positive, but flows from the periphery into the core had mixed signs. To see if this is what is going on, the next four blocks of estimates further drill down into the combinations of core and periphery trade relationships. The results here at least broadly back up these predictions, as the gap between Core-to-EMU and Periphery-to-EMU appears to be driven entirely by asymmetries across these lines. I find similar magnitude of trade effects for core-to-core, periphery-to-periphery, and core-to-periphery, with only weak evidence supporting the strong improvement of trade between core countries. Conversely I find weak evidence suggesting a roughly 7% decline in exports from the periphery into the core.

Broadly these estimates are consistent with what is implied in the conceptual framework of [section 2](#), but I note that the magnitudes are not entirely in line with what one might expect from that section, particularly given the motivating evidence I presented. From the pure theory one might have expected flows from core exporters to the periphery to be of larger magnitude than other groups. This is because these exporters should benefit both from reduction of exchange volatility, which is the largest pre-EMU volatility group in [Figure 1](#), but that they should also see relative exports improve if other forms of price/cost convergence which work to make their exports more competitive. Within core exports should see a positive trade effect, but given the relatively small size of exchange volatility pre-EMU one might have expected this magnitude to be noticeably below the other three. While the core-to-periphery estimate is among the robust estimates of a positive euro effect, it has a magnitude that is smallest (if statistically indistinguishable) of the three. One explanation that could both rationalize this and the perhaps surprisingly large point estimate of within core exports is that core exporters saw extremely strong convergence of export prices over this period, which can be seen from [Figure 2a](#) and [Figure 2b](#) which appear to converge strongly relative to the periphery relative exports shown in [Figure 2c](#) and [Figure 2d](#). If core integration is stronger than periphery it might explain larger trade

effects. To the degree that some¹⁴ of this convergence can be attributed to the euro this would reflect a true positive impact, though it is possible these estimates are also picking up some variation related to further euro integration that is not captured by regional trade deals and fixed effects.

5. CONCLUSIONS

I find empirical evidence that the euro effect on trade is not the same for core and periphery trading partners. Estimates of exports from the core suggest a 12-14% improvement in trade, which is much larger than the consensus implied by existing work on an aggregate euro effect on trade. This estimated bump in exports is similar in size whether these flows are to other core destinations or to the periphery, though statistically the estimates of the latter are much more robust. The periphery provides a different story, with zero trade effect of the EMU generally, which masks a robust improvement of trade, on the order of 13% when exporting to other members of the periphery, and weak evidence of a negative euro effect on exports into the core. These results are consistent with a framework where the introduction of the euro not only removed exchange rate volatility, but also led to price convergence core and periphery countries where the European periphery appreciated relative to the core.

On their face these results may seem pessimistic. The much vaunted trade integration resulting from currency union membership has been seen as one of the key benefits that member states can expect, particularly in the face of challenges in balancing needs across diverse member states that have arisen through the global financial crisis. However, if the mechanism driving a lack of trade improvement among peripheral members is due to economic convergence toward the core then it is possible that non-trade benefits dwarf any concerns over a missing euro-trade effect. The simple conceptual framework laid out in this paper gives credence to the idea that this is plausible, but requires deeper investigation to fully understand. More work needs to be done to flesh out the mechanism driving these results, and to systematically document the ways in which a joint euro can lead to such integration. For now we can say that from the perspective of trade there still appear to be two eurozones: a core that enjoys clear benefits from union membership, and a periphery that does not.

¹⁴I wish to be careful to note that other contemporaneous policy likely promoted this convergence, in addition to any that comes through the EMU itself.

During the preparation of this work the author used GPT-5 in order to provide assistance in: code debugging; checks for spelling, grammar, consistency of notation; and to provide feedback and quality checks on the author's own work and reasoning related to literature and methodological choices. While this tool proved useful in many respects, I would compare its work to that of a quality undergraduate RA and any output was treated with similar scrutiny. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article. Any errors are my own.

REFERENCES

- Abadie, A., S. Athey, G. W. Imbens, and J. M. Wooldridge (2023). When should you adjust standard errors for clustering? *The Quarterly Journal of Economics* 138(1), 1–35.
- Arestis, P. and M. Sawyer (2011). The design faults of the economic and monetary union. *Journal of Contemporary European Studies* 19(1), 21–32.
- Atoyan, R., J. Manning, and J. Rahman (2013). Rebalancing: Evidence from current account adjustment in europe. Technical report, IMF Working Paper WP/13/74.
- Badinger, H. and F. Breuss (2006). Country size and the gains from trade bloc enlargement: an empirical assessment for the european community. *Review of International Economics* 14(4), 615.
- Badinger, H. and F. Breuss (2009). Country size and the trade effects of the euro. *Review of World Economics* 145(2), 207–223.
- Baldwin, R., F. Skudelny, and D. Taglioni (2005). Trade effects of the euro: evidence from sectoral data. *ECB Working Paper No. 446*.
- Baldwin, R. and D. Taglioni (2007). Trade effects of the euro: A comparison of estimators. *Journal of economic integration*, 780–818.
- Bayoumi, T. and B. Eichengreen (1992a). Shocking aspects of european monetary unification. *NBER Working Paper No. 3949*.
- Bayoumi, T. and B. Eichengreen (1992b). Shocking aspects of european monetary unification.
- Belke, A., C. Domnick, and D. Gros (2017). Business cycle synchronization in the emu: Core vs. periphery. *Open Economies Review* 28(5), 863–892.
- Campbell, D. L. (2013). Estimating the impact of currency unions on trade: solving the glick and rose puzzle. *The World Economy* 36(10), 1278–1293.
- Campbell, D. L. and A. Chentsov (2023). Breaking badly: The currency union effect on trade. *Journal of International Money and Finance* 136, 102871.
- Campos, N. F. and C. Macchiarelli (2016). Core and periphery in the european monetary union: Bayoumi and eichengreen 25 years later. *Economics Letters* 147, 127–130.
- Chintrakarn, P. (2008). Estimating the euro effects on trade with propensity score matching. *Review of International Economics* 16(1), 186–198.
- Conte, M., P. Cotterlaz, and T. Mayer (2021). The cepii gravity database. *CEPII: Paris, France*.

- De Santis, R. and T. Cesaroni (2016). Current account 'core-periphery dualism' in the emu. *The World Economy* 39(10), 1514–1538.
- Feenstra, R. C., R. Inklaar, and M. P. Timmer (2015). The next generation of the penn world table. *American economic review* 105(10), 3150–3182.
- Fernández-Villaverde, J., L. Garicano, and T. Santos (2013). Fiscal unions redux. *Review of Economic Studies* 80(4), 1338–1389.
- Flam, H. and H. Nordström (2007). Explaining large euro effects on trade: the extensive margin and vertical specialization. *Scandinavian Journal of Economics* 109(4), 791–816.
- Glick, R. and A. K. Rose (2016). Currency unions and trade: A post-emu reassessment. *European Economic Review* 87, 78–91.
- Head, K. and T. Mayer (2014). Gravity equations: Workhorse, toolkit, and cookbook. In *Handbook of international economics*, Volume 4, pp. 131–195. Elsevier.
- Hirano, K. and G. W. Imbens (2001). Estimation of causal effects using propensity score weighting: An application to data on right heart catheterization. *Health Services and Outcomes research methodology* 2(3), 259–278.
- Imbens, G. W. (2004). Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and statistics* 86(1), 4–29.
- Jordà, Ò. and A. M. Taylor (2016). The time for austerity: estimating the average treatment effect of fiscal policy. *The Economic Journal* 126(590), 219–255.
- Kenen, P. B. (2002). Currency unions and trade: Variations on themes by rose and persson. *Reserve Bank of New Zealand Discussion Paper No. DP2002/08*.
- Kopeccky, J. (2023). Many unions, one estimate? disaggregating the currency union effect on trade. *Emerging Markets Finance and Trade*, 1–23.
- Kopeccky, J. (2024, February). A match made in maastricht: Estimating the treatment effect of the euro on trade. *Open Economies Review* 35(1), 159–187.
- Lane, P. R. and B. Pels (2012). Current account imbalances in europe. *Monetary Authority of Finland Research Discussion Papers* (5/2012).
- Lunceford, J. K. and M. Davidian (2004). Stratification and weighting via the propensity score in estimation of causal treatment effects: a comparative study. *Statistics in medicine* 23(19), 2937–2960.
- McKenzie, D. (2017). When should you cluster standard errors? new wisdom from the econometrics oracle. *World Bank Blogs*.
- Mink, M., J. Jacobs, and J. de Haan (2007). Measuring synchronicity and co-movement of business cycles within the euro area. *Journal of Business Cycle Measurement and Analysis* 3(1), 1–27.
- Montoya, L. A. and J. de Haan (2008). Regional business cycle synchronization in europe? *International Economics and Economic Policy* 5(1), 123–137.
- Persson, T. (2001). Currency unions and trade: how large is the treatment effect? *Economic policy* 16(33), 434–448.

- Polák, P. (2019). The euro's trade effect: a meta-analysis. *Journal of Economic Surveys* 33(1), 101–124.
- Puzzello, L. and P. Gomis-Porqueras (2018). Winners and losers from the euro. *European Economic Review* 108, 129–152.
- Rose, A. K. (2000). One money, one market: the effect of common currencies on trade. *Economic policy* 15(30), 08–45.
- Rose, A. K. (2001). Currency unions and trade: the effect is large. *Economic Policy* 16(33), 449–461.
- Rose, A. K. (2017). Why do estimates of the emu effect on trade vary so much? *Open Economies Review* 28(1), 1–18.
- Silva, J. S. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and statistics* 88(4), 641–658.
- Silva, J. S. and S. Tenreyro (2011). Further simulation evidence on the performance of the poisson pseudo-maximum likelihood estimator. *Economics Letters* 112(2), 220–222.
- Słoczyński, T. and J. M. Wooldridge (2018). A general double robustness result for estimating average treatment effects. *Econometric Theory* 34(1), 112–133.
- Wooldridge, J. M. (2007). Inverse probability weighted estimation for general missing data problems. *Journal of econometrics* 141(2), 1281–1301.