Why estimates of the trade effects of the Eurozone vary so much

Andrew Rose

19 October 2016

The pro-trade effects of the euro are a clear-cut benefit of Eurozone membership, but scholarly estimates of the size of this effect vary widely. This column uses meta-analysis to argue that the variation stems from inappropriate exclusion of nations and years. When all countries and years available in the data are included, the estimate of the euro trade effect is economically and statistically large, at about 50%.

When scientists want to summarise an entire literature quantitatively, they use the tools of meta-analysis. This involves combining together all the estimates of a particular coefficient using standard statistical techniques; my interest in this column is the effect of the euro on international trade. Luckily, researchers use essentially the same methodology to estimate this coefficient. They link the amount of trade between, say, France and Germany to the distance between France and Germany and the economic masses (think GDP) of France and Germany, along with a few other statistical bells and whistles. This 'gravity' model of trade is one of the most successful empirical models operating in any field of economics, as surveyed by Head and Mayer (2014). Once the gravity model is used to explain most of the variation in trade, it becomes easy to estimate if there is any leftover effect for countries that share a common currency.

Related
The currency union effect on trade: Redux
Reuven Glick, Andrew Rose
How did the euro boost trade? Microeconomic evidence
Volker Nitsch, Mauro Pisu

There are 45 papers circulating that estimate the effect of the euro on trade (only 25 of these are published). From each of these studies I collect the authors' preferred estimate of the effect of the euro on trade, along with its standard error and other study features.

Figure 1 is a 'Forest plot' that summarises the literature. The (45) papers are listed on the left; separate columns list when the data used in the study begin and end, and how many countries are included in the study. Three other technical features are also tabulated on the left: a) whether the study examines bilateral exports or trade; b) whether it includes dyadic fixed effects; and c) whether it includes time-varying country and fixed effects. Since the work of Baldwin and Taglioni (2007), there has been a growing consensus among the (admittedly nerdy) cognoscenti that reliable studies should examine exports (not two-way trade) and include both sets of fixed effect. So, the more of these features present in a study, the more trustworthy the results. The study effect – the trade effect of the euro – is tabulated on the right, along with a 95% confidence interval.

Figure 1 Forest plot of 45 literature estimates of the euro effect on trade/exports

Andrew Rose
Professor of International Business, in the Economic Analysis and Policy Group, Haas School of Business, University of California, Berkeley; and Research Fellow, CEPR

Don't Miss
The bond-equity allocation of the Norwegian sovereign wealth fund
Henriksen, Mork
Using bibliometrics to gauge research quality
Régibeau, Rockett
Hysteresis and fiscal policy during the Global Crisis
Fatás, Summers
Time for growth
Boerner, Severgnini
Heterogeneity and unemployment dynamics
Ahn, Hamilton
How not to build a state: Evidence from Colombia
Acemoglu, Fergusson, Robinson, Romero, Vargas
What big data tells us about real income growth
Redding, Weinstein
Evaluating Trump's trade policies
Hufbauer, Jung
The objectives of financial stability policy
Tucker
Conceptual challenges in international finance
Avdjiev, McCauley, Shin
Happiness inequality and the importance of trust
Helliwell

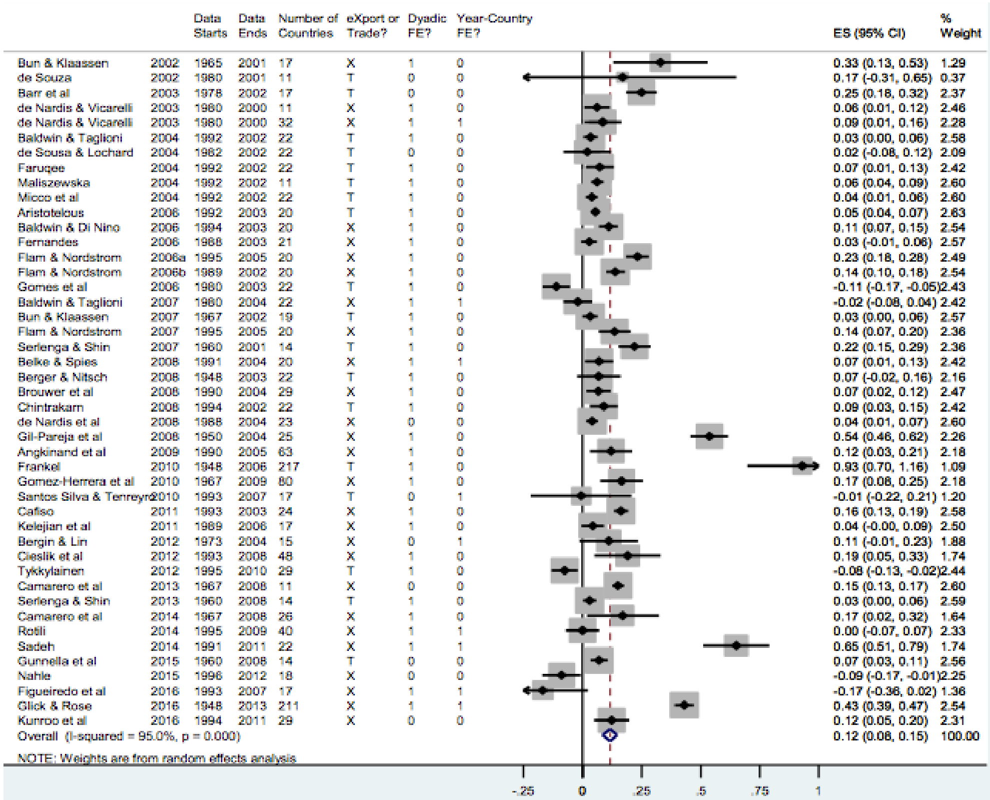

How does the literature collectively quantify the euro trade effect? The standard meta-analysis 'random effects' estimator (the vertical line in Figure 1 marked with a diamond in the bottom row) indicates a statistically and economically significant euro trade effect of (exp(.116)-1≈) 12.3%. This effect is significantly different from zero on both economic and statistical terms. Nothing to sneer at, but … also nothing to write home about.

Why do estimates of the euro trade effect vary so much?

The most striking thing about the forest plot of Figure 1 of is how much the estimates vary across studies. A bunch are enormous, exceeding .5, while six of the study effects are actually negative. Which leads me to the titular question of this column: Why do the estimates vary so much? It turns out that there are two reasons: the number of countries included in the data set and the number of years. Including either more countries or more years (or both) in the data set leads to higher estimates of the effect of the euro on trade.

It is easy to show that the euro trade effect on trade rises systematically when more countries and more years are included. I show this in Figure 2, which presents estimates of the euro effect on exports from my recent paper (Rose 2016), along with a +/- two standard error confidence interval. I use a standard plain-vanilla gravity model of exports between countries, and only vary the sample of countries and years.

The top line of Figure 2 connects estimates of the euro export effect that use data for all (>200) countries. I begin the sample in 1948 and initially end it in 2001, two years after the euro technically began and the year before the it was physically introduced. I then add more data, in two-year increments. The middle line is similar, but includes only data for rich countries (using the World Bank definition of annual GDP per capita of at least $12,736). I have data on >34,000 country-pairs, but there are only ≈4,000 pairs of rich countries, an order of magnitude lower. The bottom line is an analogue for the approximately 800 country-pairs ever inside the EU (now some 28 countries, with 19 in the Eurozone).

Figure 2 Estimates of the euro effect on trade with varying data samples

What else can central banks do?
Ball, Gagnon, Honohan, Krogstrup

Brexit: Economic prosperity and voting behaviour
Zoega

A new eBook: Brexit Beckons
Baldwin

How secular stagnation spreads and how it can be cured
Eggertsson, Summers

The greatest reshuffle of individual incomes since the Industrial Revolution
Milanovic

Making the Eurozone more resilient: What is needed now and what can wait?
Resiliency Authors

New VoxEU eBook: How to fix the Eurozone
Baldwin, Giavazzi

Most Read

This Month | All Time

Piketty's housing capital results: New US facts
La Cava

How computer automation affects occupations
Bessen

Why a banking crisis in China seems unavoidable
Campanella, Vernazza

Income inequality in a globalising world
Niño-Zarazúa, Roope, Tarp

Democracy does not cause growth
Ruiz Pozuelo, Slipowitz, Vuletin

more

CEPR Policy Research

Discussion Papers | Insights

Homeownership of immigrants in France: selection effects related to international migration flows
Gobillon, Solignac

Climate Change and Long-Run Discount Rates: Evidence from Real Estate
Giglio, Maggiori, Stroebel, Weber

The Permanent Effects of Fiscal Consolidations
Summers, Fatás

Demographics and the Secular Stagnation Hypothesis in Europe
Favero, Galasso

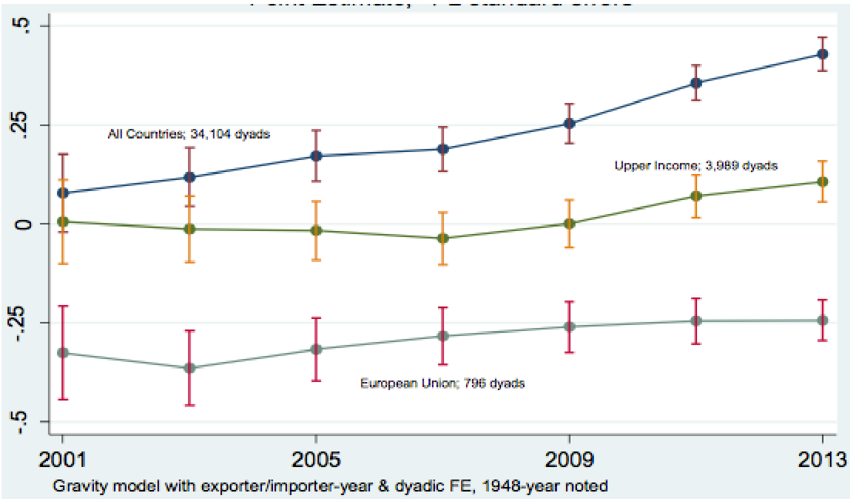

There are two interesting findings in Figure 2. First, the estimates tend to rise with the sample's span of time. Second, euro estimates that employ data from the entire world are positive and economically significant, the upper-income analogues are usually insignificantly different from zero, and the EU data deliver significantly negative point estimates of the euro export effect, despite the latter's understandable focus on the data most obviously relevant to the euro.

All this is reassuringly consistent with the literature summarised in the Forest plot of Figure 1, and essentially fills in gaps in the literature. It points to three conclusions. First, throwing away data – as much of the literature does – easily allows one to estimate a small and/or negative export effect of the euro. Second, increasing the span of time by adding observations relevant to the euro tends to increase the estimated euro export effect. Finally, increasing the span of countries – adding observations for countries that seem irrelevant to the euro – also increases the euro export effect. While the first two observations are intuitive and appealing, the last seems odd.

Why do extra data matter?

The evidence from both the literature/meta-analysis and empirical work is clear. The estimated euro export effect rises as the cross-section of countries expands, even if the extra countries appear unconnected to the euro. Why?

The answer is mostly technical. A revolutionary development in the gravity model, discovered by Anderson and van Wincoop (2003), forced researchers to consider the importance of "multilateral indices of 'trade resistance'", which changes the way gravity models need to be estimated. Technically speaking, one should always include "time-varying exporter and importer fixed effects". Dropping small and poor countries leads to biased estimates of these effects if small and poor countries have systematically different 'trade resistance', or if the trade resistance of large countries is reflected in trade with smaller countries. Intuitively, the fact that Germany exports to small countries like Fiji is important in understanding German export prowess, but this information is lost if Fiji isn't included in the sample of data. This bias helps explain why estimates of the euro trade effect change systematically with the breadth of the data sample – the estimated trade effect is biased downward when small/poor countries are omitted from the sample and the fixed effects change accordingly. As multilateral trade resistance is a function of all bilateral trade barriers, all trade partners should be included for the most accurate estimates.

So where do we end up? When you include all countries and years available in the data set – as seems wise – the estimate of the euro trade effect is economically and statistically large. My data includes more than 200 countries and ends in 2013; by that point, the point estimate of .43 corresponds to a substantive economic impact of (exp(.43)-1≈) 54%, with a robust t-statistic of over 20. Both statistically and economically, the effect is large.

Where do we stand?

Why do estimates of the euro trade effect vary so much? Including more observations – over time, or especially by country – increases the estimated effect of the euro on trade. In recent work, I have established this stylised fact with meta-analysis and empirics, and then interpreted it. The dependency of the euro export effect on the number of countries is more striking and less intuitive than the time-dependency. My explanation is that truncating the sample by omitting small/poor countries biases downward the estimates of the country-time fixed effects. This leads to downward bias in the estimated partial effect of the euro on exports. Succinctly, one can shrink the estimated euro export effect by inappropriately dropping observations. In general, there is little reason to drop data without necessity; in this case, it is important to include as much data as possible when estimating the euro trade effect.

References

Anderson, J E and E Van Wincoop (2003), "Gravity with Gravitas: A Solution to the Border Puzzle", *American Economic Review* 93(1), 170-192.

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

Head, K and T Mayer (2014), "Gravity Equations: Workhorse, Toolkit, and Cookbook", Chapter 3 in G Gopinath, E Helpman and K Rogoff (eds), *Handbook of International Economics*, Vol. 4, Amsterdam: Elsevier, pp. 131–195.

Rose, A K (2016), "Why Do Estimates of the EMU Effect on Trade Vary So Much?", CEPR Discussion Paper No. 11532.

Topics: International trade

Tags: euro trade effect, eurozone, meta-analysis, estimates

Related

The currency union effect on trade: Redux
Reuven Glick, Andrew Rose

How did the euro boost trade? Microeconomic evidence
Volker Nitsch, Mauro Pisu