UCLouvain

UNIVERSITÉ DE NAMUR

Economics School of Louvain - ESL

Economics School of Namur - ESN

# The Skill Content of Value-Added Exports and Populist Voting Behavior

Author : Federico J.F.J. Willems

Thesis Director : Gonzague Vannoorenberghe

Academic Year 2025-2026

Master in Economics – 60 credits

Université catholique de Louvain - Economics School of Louvain - Place Montesquieu 3, 1348 Louvain-La-Neuve
Université de Namur – Economics School of Namur – Rempart de la Vierge 8, 5000 Namur

# Abstract

Gross trade statistics record where goods cross borders, not where value is created so they miss the linkages that propagate through the value chain. This paper examines one such linkage: the foreign value added share of exports and its relationship with populist voting. We measure populism by the volume margin of Docquier et al. (2026). Because a large part of what any country exports is already embedded in the chain, removing the double counting in gross trade statistics changes which countries look exposed to globalisation.

Estimating by Poisson pseudo-maximum likelihood with country and year fixed effects, a one percentage point widening of the gap between low-skill and high-skill foreign value added share is associated with an increase of approximately 10.6% in the populist volume margin ($\beta = 0.101$, SE 0.049, $p = 0.040$; 269 observations across 39 countries). The effect is concentrated entirely in left-wing populism ($\beta = 0.303$, SE 0.088, $p < 0.01$), with no detectable right-wing response. The asymmetry is what Rodrik's (2021) framework predicts rather than an anomaly. Production fragmentation is a trade and investment shock and such shocks activate the income cleavage and mobilise the left. The ethno-national cleavage that mobilises the right is activated by immigration. Significance holds when sector shares are weighted by export value, though the coefficient falls by roughly half.

Whether exposure is measured by simple mean or by export weight therefore determines what one concludes about who globalisation harms. Offshoring and import competition activate different politics and treating them as interchangeable obscures a distinction voters appear to be making.

# Table of Contents

# 1. Introduction

Gross trade statistics record where goods cross borders, not where value is created. This distinction matters because it changes which countries look exposed to globalisation. Koopman, Wang and Wei (2014) show how gross exports can be decomposed into their value-added origins, removing the double counting that arises when intermediate goods cross borders repeatedly.

The rise of populism across advanced economies is one of the defining political trends of the last two decades. The literature demonstrates that imports of low-skill intensive goods reduce local manufacturing employment, leaving workers displaced and economically insecure. Populists take advantage of this insecurity, offering simple and powerful rhetoric in answer to complex problems. Because gross flows attribute a good's entire value to its final exporter, any measure of the skill content of trade inherits that misattribution.

Populism resists easy definition. Guriev and Papaioannou (2022) warn against the temptation to apply the label to any movement one happens to dislike, and Mudde (2004) argues that the concept is thin-centred. It has a minimal core and attaches itself to a host ideology which is why populist parties of the left and the right can look so different while sharing the same underlying structure. Any empirical treatment therefore needs an operational definition rather than a judgement call. This paper follows Docquier et al. (2026) in characterizing a populist party by two dimensions of its manifesto, its anti-establishment stance and its commitment to economic protection.

This paper contributes to the literature by examining an alternative channel of trade exposure: trade in value added. Building on Docquier et al. (2026) and using the same underlying Manifesto Project data, I test whether the finding that low-skill imports increase populism extends to the foreign content embedded in exports. Replacing their gross import measure with the foreign value added share of gross exports (FVASH) from the OECD TiVA database eliminates the double counting inherent in raw trade statistics. FVASH captures backward GVC participation, measuring the share of a country’s exports originating in foreign inputs. Manufacturing sectors are disaggregated into low-skill and high-skill groups. This paper therefore asks whether the foreign value added content of manufacturing exports, disaggregated by skill intensity, predicts populist voting across OECD countries.

Four strands of literature meet in this question. The first establishes that gross trade statistics misattribute value across countries: Koopman, Wang and Wei (2014) decompose gross exports into their value-added origins, and Johnson and Noguera (2012) show that intermediates account for as much as two thirds of world trade. This work stops at the level of countries and industries and does not connect the correction to labour markets or to political outcomes.

The second strand does connect trade to politics. Colantone and Stanig (2018) find that regional import exposure raises support for nationalist parties across fifteen Western European countries, and that the effect extends beyond the directly displaced to service and public-sector workers in affected regions. Autor et al. (2020) document political polarisation following Chinese import competition in the United States. Docquier et al. (2026) show that imports intensive in low-skill labour raise populist vote shares across a panel of 55 countries. What unites these papers is that exposure is measured in gross terms.

That is the gap this paper addresses. Docquier et al. measure the skill content of gross imports, but gross imports attribute a good's entire value to its last exporter. The skill content of a gross flow is contaminated by value added originating elsewhere. This paper measures the skill content of trade in value-added terms instead.

A second contribution is methodological. Sector-level shares must be aggregated to a group before they can be used as a regressor, and the choice of aggregation rule is not innocuous. Sector shares can be aggregated with equal weight or by export value, and under the two rules a different skill group appears more foreign-sourced.

The main results are as follows. A one percentage point widening of the gap between low-skill and high-skill foreign value added share is associated with an increase of roughly 10.6% in the populist volume margin, estimated by Poisson pseudo-maximum likelihood with country and year fixed effects. The effect is concentrated entirely in left-wing populism and is statistically indistinguishable from zero for right-wing populism.

# 2. Motivation

## 2.1 Issues with gross trade statistics

A gross export figure records who sold a good, not who made it. When we take a gross trade figure, we do not care at what part of the global valued chain the good was taken or whether it was a final good or intermediary good as long as it was exported. This causes the export numbers to be inflated as your export industry workers or industry is not fully benefiting from the full value of the good if it at a later stage of production.

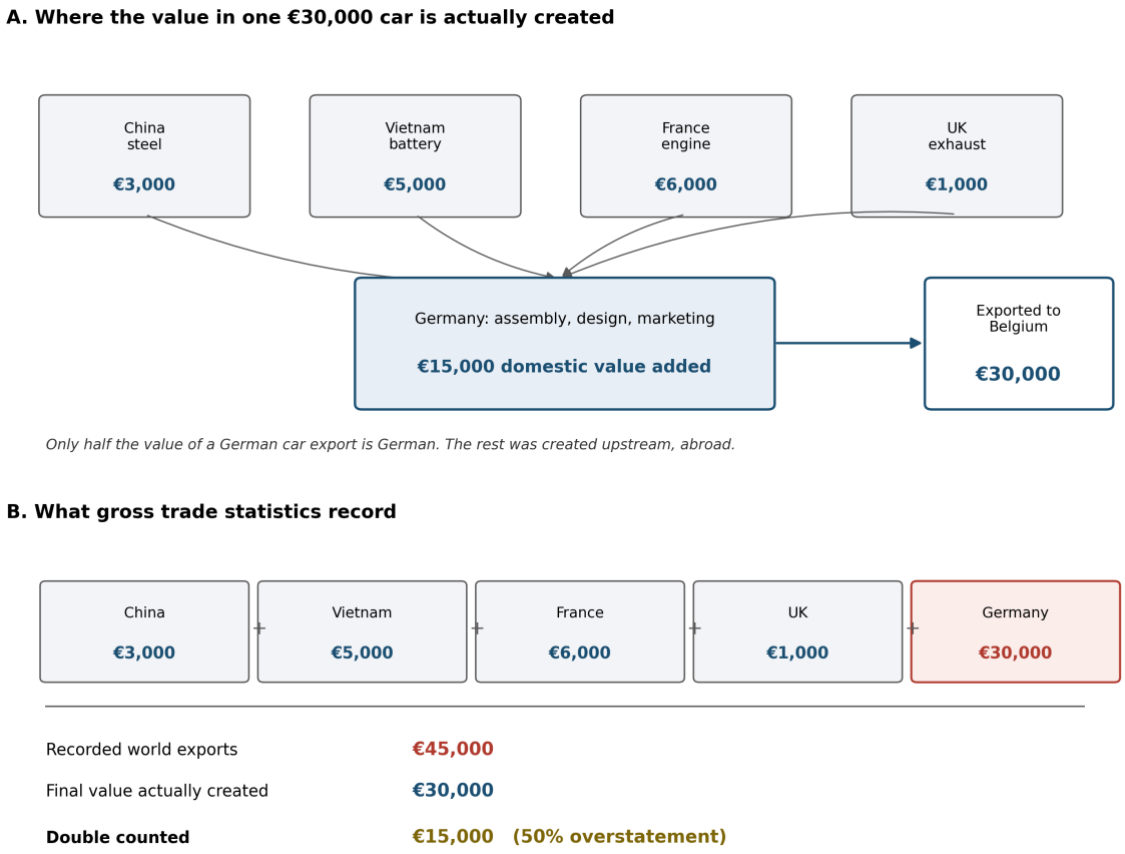


**Figure 1. Double counting in a single automotive value chain.** Illustrative figures. A car assembled in Germany and sold in Belgium for €30,000 contains components made in four other countries. Each component is recorded as an export by its producer, and the finished car is recorded again at its full value, so world exports of €45,000 are registered against €30,000 of value created.

Let's take a simple example. A car assembled in Germany and sold in Belgium for €30,000 contains components made in four other countries. Each component crosses a border and is recorded as an export by its producer. The finished car is then recorded again at its full value. World exports of €45,000 are recorded against €30,000 of value created a 50% overstatement. When Audi exports a car to the Belgian market, the exports will report the final value of that car. But when we zoom in and look at the stages of production of a car, the steel may have been imported from China, the batteries from Vietnam, the engine from France, the exhaust from Britain and so on. This amounts to a value added for this car being created in all these different countries and the car was only produced in Germany. Not only do the German final export include this value added in which German workers do not benefit, they are still

counted by the Chinese market, Vietnamese market, British and French markets too. This inflates the true value that exports have to a country. The global valued chain is affected by all sorts of trade and business strategies hence when we look at trade, we must look at the GVC or our conclusions may be a false representation of reality In fact, according to Johnson & Noguera (2012) intermediates are up to two-thirds of world trade. This is a vital part gross statistics miss and the reason for it can be a whole plethora of reasons business based[1]and country based[2] reasons.

Zooming into value added literature, one part has tried to explore the valued added channels and tried to explore the decomposition of gross statistics in papers such as Koopman, Wang and Wei (2014) for the decomposition framework, where the theoretical foundation of value added comes from. Starting from gross exports they derive an input-output decomposition into nine components. Eliminating double counting can therefore reverse standard trade conclusions: the paper demonstrates that the US when assessed using value added rather than gross trade moves from 10th to 3rd in revealed comparative advantage in metal products switching from a comparative disadvantage to a comparative advantage. A similar input-output framework underpins the OECD TiVA database which is the primary data source for this thesis. Johnson and Noguera (2012) also provide the distinction between double counting and indirect routing through third countries.

Hence, when we look beyond the naïve statistics, we discover a whole heap of conflicting channels in the global valued chain leading to the demise of manufacturing globally in Western markets. if gross flows misattribute value across countries, then any measure of the skill content of gross flows is contaminated. The skill content of a good is a claim about who performed the work. This is the precise sense in which the paper expands on Docquier et al. (2026).

## 2.2 Why exports, and why the foreign value-added share

When exploring valued added, we need to make an executive decision on what we would like to explore and its consequences for our conclusions. The OCED VA trade statistics cover both forward and backwards participation. VA can be explored through decomposing exports, imports, final demand, levels or shares. [3]

It is important to note we have limited the sample to only manufactured goods due to the accessibility of data and how services, or other local production is convoluted when looking at a valued chain preceptive. By sticking to one specific channel, manufacturing, we can provide coherent economic reasoning for it.

For example, agriculture was rejected as it usually only has one or two production stages, is largely non-traded, is dominated by domestic consumption and by subsidy regimes that vary enormously across countries. Cross-country FVASH comparisons could reflect policy more than fragmentation. Take services for example too. Services are intangible, poorly measured in input-output tables, and often delivered without crossing a border in any recorded way,

[1] for example, for FDI reasons, access to markets, cheaper transport, exploit lower variable costs etc

[2] For example, free trade agreements, tariffs, NTBs, geography, political affiliation etc

[3] Backward participation measures foreign value added embedded in a country's own exports, so the country is the buyer of foreign inputs. Forward participation measures domestic value added embedded in other countries' exports, so the country is the upstream supplier. Only the backward measure captures the relocation of production stages away from the domestic economy.

which is precisely why TiVA coverage of services value added is weakest. Manufacturing is where fragmentation is both economically meaningful and statistically visible.

Due to the vast nature of value added, it is important to state why I have rejected some of the alternative measures to exports. When looking at imports the mechanism would work through measuring foreign competition. Foreign goods arriving and displacing domestic products. This literature has already been covered extensively by import penetration and papers such as Autor et al. (2020) and Colantone and Stanig (2018). It links clearly with populism by Rodrik (2018). Rodrik notes estimates attributing only 10–13% of US manufacturing job losses in the 2000s to import competition. The perceived fairness of a shock matters independently of its size, which is how a channel responsible for a minority of job losses can drive most of the political reaction. This exposes the industries that are worse effected in which local firms shut down as are not competitive enough displacing workers in these industries.

Value added (VA) in exports, on the other hand, measure the composition of your own production, how much you sell abroad that was produced by domestic workers. Therefore, any data that exists is from an industry that has survived foreign competition and is competitive enough to export. Import competition kills the firm but offshoring hollows out its low-skill stages while the firm survives and keeps exporting. Therefore, it is also important to investigate exports even if it is a novelty and comes with its challenges. In addition, exports are directly comparable and have great data as every time they cross a border, they are tracked by national authorities. Exports are measured consistently across all 45 countries in the same currency and the same year, so no deflation or exchange-rate conversion is needed. Secondly, as exports are at the border, offshoring is traceable and visible as the relocation registers in trade flows by construction.[4]

Exports are not perfect either. The measure does not capture offshoring by firms serving only the domestic market as the final product will never cross the border. It is less informative for economies with small export sectors in which the offshoring story may be less prevalent when evaluating the increase of populism.

The data will be interpreted in shares. The share is a ratio of two contemporaneous flows in the same currency, so the price level cancels. It needs no deflation and is immune to revaluation. This is useful when all countries and industries vary heavily in size.

One feature of the panel should be stated at the outset. The 45 countries covered by TiVA are predominantly high-income, since coverage follows OECD membership and its principal partners. This suits the mechanism under test, which is defined for high-wage economies relocating labour-intensive stages abroad, but it bounds what the results can speak to.

## 2.3 Exploring FVASH

When exploring VA exports, it is important to directly compare it to gross exports to show how important double counting inflates the figures. To do this I compared the trends of gross exports of the 45 countries in my sample compared to VA in exports and VA in foreign domestic exports.

[4] VA denotes value added throughout.

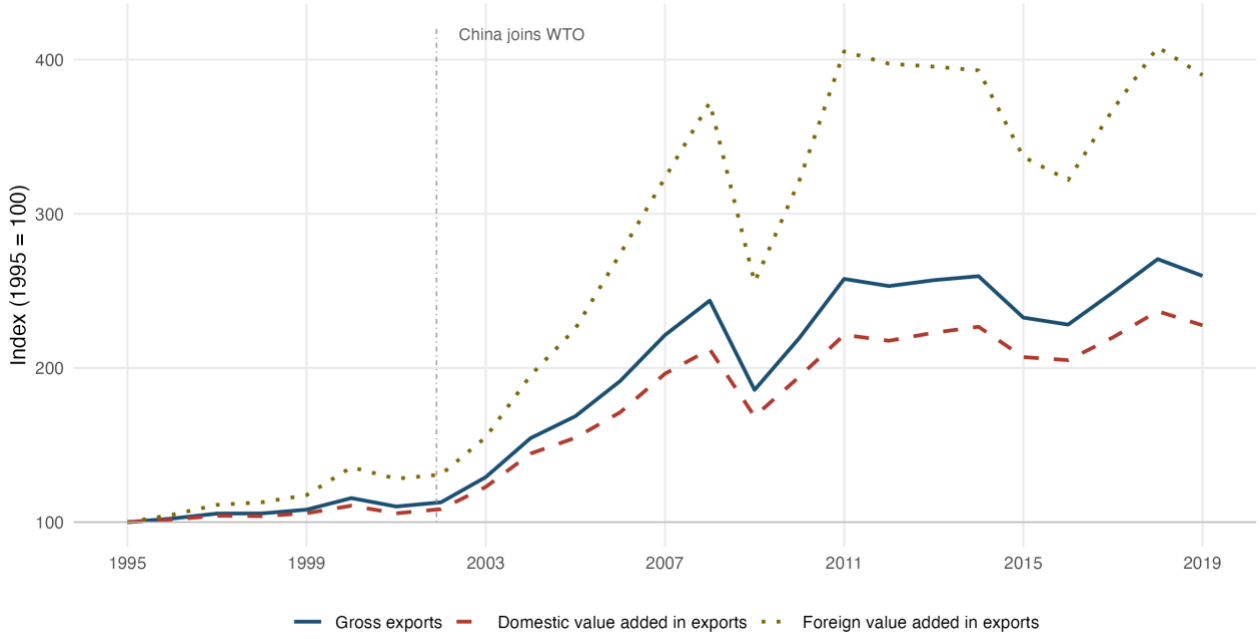


**Figure 2. Gross exports and their value added components, indexed to 1995.** Manufacturing aggregate (ISIC Rev. 4), pooled across the 45 countries in the panel. Series are pooled across countries, Source: OECD TiVA database (2025 edition).

The graph supports the fact we should look at gross exports with a pinch of salt. Before 2002, all 3 factors have similar trends. Over the period, when indexing to 100 in 1995, we find gross exports to rise over time with a sudden dip at the financial crisis. Value added in domestic exports follows a similar trend but its gap with exports widens from 2004 suggesting a larger part is due to foreign valued added. In fact, when we look at foreign value added it has effectively quadrupled over the time with a similar dip in the financial crisis. It is important to note that these are indexed figures so to start with the foreign value added would have been much lower than its counterpart so quadrupling a small number many are not as significant as this graph may show it. What is interesting is that it mimics the structural break Docquier et al. (2026) find when they interact low-skill imports with a post-2001 dummy. Also note the 2009 collapse is asymmetric, foreign value added falls by roughly 30% against 20% for domestic. Henceforth, gross statistics mistake volatility as well as levels. When firms are hit by volatility, they restrict exports and focus on domestic production (Bems, Johnson and Yi, 2013). The graph shows compelling evidence for why we should explore the decomposition of exports and that its interpretation for how it effects labour show be nuanced as export benefits are disproportional distributed.

The way we would link it to populism is through the following. Higher foreign value added share means a larger share of production stages performed abroad. The domestic value added share falls by the identity.[5] The stages that relocate are disproportionately low-skill intensive as companies have an incentive to outsource to foreign countries with lower variable costs. Higher skilled manufacturers are still paid a decent wage and have more specialised training that are hard to replace therefore you cannot substitute with a random worker abroad. A process engineer or toolmaker embodies firm-specific knowledge that takes years to rebuild elsewhere, an operator on a standardised line does not. The offshoring results in displaced

[5] explained in methodology but 100% = DVASH share + FVASH share

workers and declining industrial regions generate economic grievances[6]. The grievance from losing work is attributed to globalisation and converts into populist voting.

This fits well with the literature. Feenstra and Hanson via Antràs, de Gortari and Itskhoki (2017) find outsourcing accounts for as much as 40% of the rise in the US skill premium. Johnson (2018) finds the Danish matched worker-firm evidence that offshoring contracts domestic employment specifically among low-skilled workers. Colantone and Stanig (2018) find falling institutional trust shifts votes. The effect is not confined to the directly exposed: service workers and public sector employees in affected regions respond too, which they read as a response to the decline of a whole industrial district rather than to individual job loss.

Colantone and Stanig's mechanism centres on the uneven distribution of globalisation gains. The 'losers' from trade support redistributive or nationalist and isolationist policies. They argue right winged populists can capture the grievances easier. Autor et al. (2020) on the other hand, find broader ideological polarisation rather than a uniform rightward shift in the US setting.

## 2.4 What foreign value added looks like across sectors

When disaggregating the manufacturing export shares, we can see a clearer picture into which sectors are more effected than others.

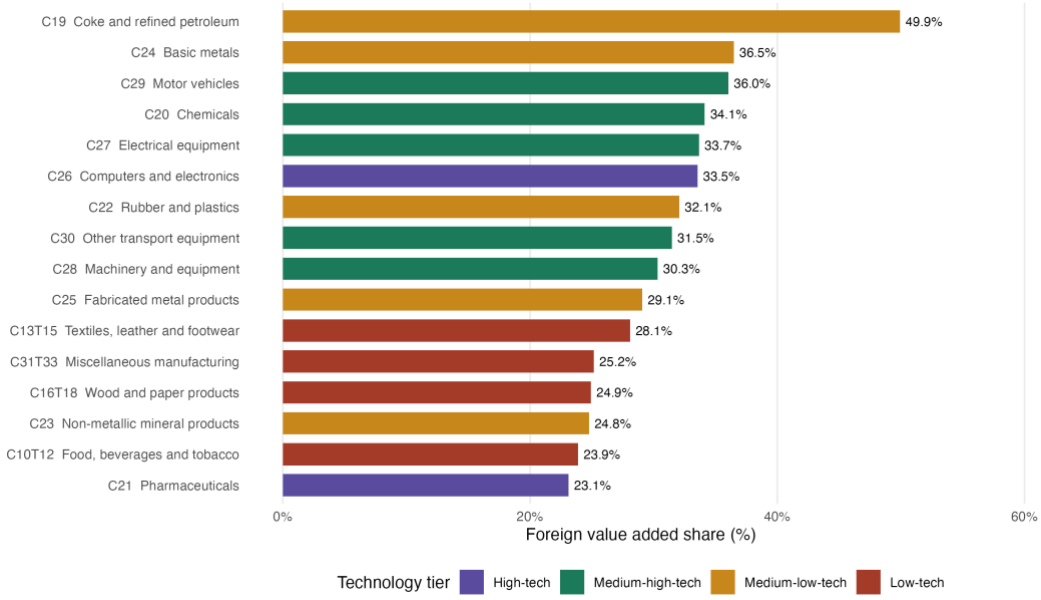


**Figure 3. Average foreign value added share by manufacturing sector, 1995–2019.** Unweighted mean across the 45 countries in the panel, excluding sector-years with no exports.Technology tiers follow Hatzichronoglou (1997). Source: OECD TiVA database (2025 edition).

From Figure 3 we see that foreign content ranges from 23.1% in pharmaceuticals meaning little to no content from abroad in manufacturing (when looking at pharmaceutical specifically, a possible economic reason is that the production is focused only in certain very specialised countries such as the US as it is mostly research rather than multiple stages of production. With coke and refined petroleum on the other hand it reaches up to 49.9% (oil has

[6] a limitation to my analysis is I do not have regional data but country level so it is impossible to know exactly what type of workers or which areas will end up voting populist, is that area a swing state? Or is it stable non populist area, attributed to regional heterogeneity

many uses in all stages of production like producing plastics, fuels and is concentrated only in few oil rich countries like the US, Iran, Russia etc). FVASH cannot on its own distinguish production fragmentation from resource dependence. There is a significant difference in VA depending on manufacturing sector. The key is the ordering does not track with technology intensity, hence there is enough within variation to argue for disaggregating. For example, the high skill sectors are both pharmaceuticals and computers and electronics, they are both neither low in FVASH or high in FVASH.

A potential issue is that the shares do not reflect how big each sector is, how much weight they should hold when constructing an accurate representation share. This can also go to the country level. If it exports even a tiny amount, the data skews for that whole sector. Additionally, the proxy for skill is technology intensity. So, when interpreting using populism, the high skill will be unclear due to the fact little people are employed in the first place. By separating into 4 sectors, we can now see how they vary over time in FVASH.

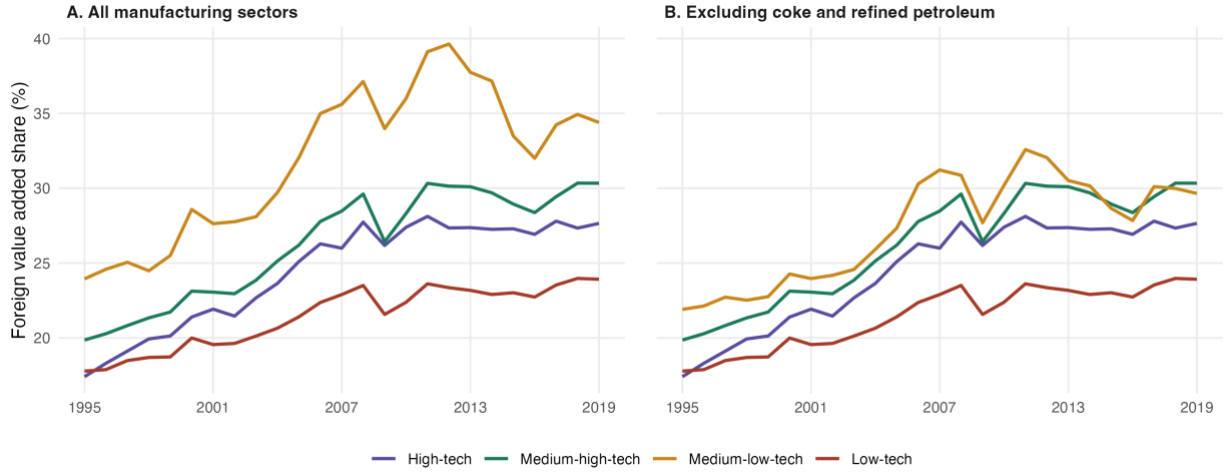


**Figure 4. Foreign value added share by technology tier, with and without coke and refined petroleum.** Export-weighted within each tier and pooled across countries. Excluding that single sector in Panel B removes most of the excess: the tier peaks at 33 per cent and converges on medium-high-tech by 2019. Source: OECD TiVA database (2025 edition).

All 4 types of manufacturing have increased over time from 1995 till 2019 with some variation in-between 2008. This diagram shows a very interesting result, theoretically we would perceive the low tech sectors in manufacturing to be the most displaced due to the Western countries (the composition of the database) having a comparative advantage in capital rather than labour intensive goods. This shows the opposite; the low sectors have been the least replaced. Some theories behind it could be the goods are bulky, often perishable, and produced close to where they are consumed. Their supply chains are short and largely national, so little foreign value enters them whatever the wage differentiates abroad. Secondly, these sectors offer little to fragment. Where production runs from raw input to finished good in one or two steps, there are few separable stages to relocate, so the option of offshoring is not really available even where it would be cheaper. Another possibility is the labour intensive low skill such as textiles have been eliminated from the survival bias of the OCED dataset, so this relationship isn't seen.

The second least displaced industry is high skilled manufacturing. Western countries have a comparative advantage in technology due to complexity of the produce and the ability to generate a larger profit margin than primary goods.

The medium tier manufacturing is the most displaced, these are more labour intensive goods that lower income countries have a comparative advantage in, from lower relative cost per worker. These sectors are integrated in the valued chain with multiple steps of production.

In Panel B I exclude coke and refined petroleum. The sector records the highest average foreign content in manufacturing, 48.8% against 36.5 for the next highest. More importantly, its foreign content means something different from the rest of the sample. Refining is capital-intensive and resource-based, and its foreign share reflects imported crude oil: a raw material most refining countries simply cannot produce rather than production stages once performed at home and since relocated abroad. FVASH cannot distinguish resource dependence from production fragmentation. Excluding coke and refined petroleum in Panel B removes most of that excess and the tiers converge. A single resource-dependent sector therefore accounts for most of the difference. We see that its trend closely follows medium high tech with FVASH being historically slightly above and now roughly the same.

Splitting the four tiers into just two skill groups tells a similar story.

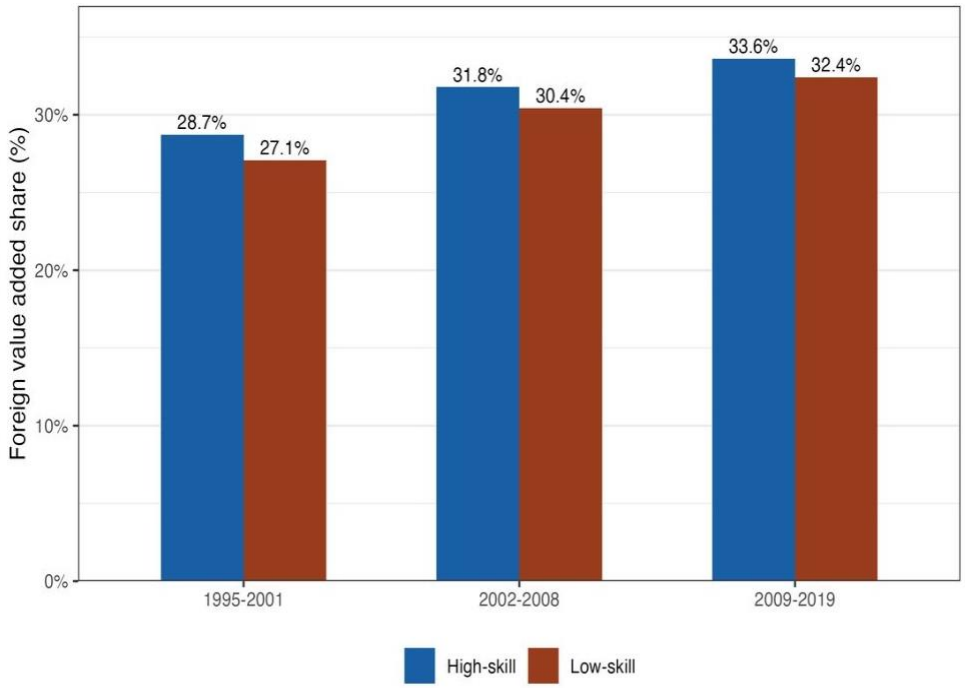


**Figure 4b. Foreign value added share by sub-period and skill group.** Unweighted mean across the 45 countries in the panel, with the four technology tiers of Figure 4 collapsed into two groups: high-skill comprises the high and medium-high tiers, low-skill the medium-low and low tiers. Source: OECD TiVA database (2025 edition).

Both groups rise across the three sub-periods: high-skill from 28.7 to 33.6%, low-skill from 27.1 to 32.4% and the gap between them narrows from 1.6 to 1.2 percentage points. The increase spans the financial crisis rather than being interrupted by it. This could suggest the sectors with the medium low and medium high hold more weigh in the aggregation (as they are closer together in Figure 4). This need not be a true representation as not weighted by sector size but still a useful insight. We see that even with the financial crisis, FVASH is increasing in all sectors over time.

This looks like a contradiction of the four-tier picture. Low-tech is the least fragmented tier; medium-low is the most. Both fall in the low-skill group, and medium-low dominates it as includes five of its nine sectors and much the larger export value. The two-way split is

therefore driven by medium-low-tech manufacturing, which is where the labour-intensive stages sit.

## 2.5 Weighted TiVA

TiVA reports FVASH by sector, but the regressor is defined for a group of sectors. Two aggregation rules are available. The simple mean gives every sector one vote regardless of size as the graph represent before. Hence, pharmaceuticals at 23% and refined petroleum at 50% count equally. The export-weighted mean takes the ratio of group totals, giving each sector influence proportional to its exports. It equals the ratio of foreign value added to gross exports for the group as a whole and reproduces the published TiVA manufacturing figure exactly. The simple mean is an average of ratios and corresponds to no observable quantity.

TiVA gives you a foreign value added share for each of 16 manufacturing sectors separately which must be aggregated. With the simple mean, each sector of the 16 has an equal weight even if some may be much larger than others level wise. Introducing the export weight solves this issue. For example, when we aggregated to 4 sectors before, it leaves two purely low skill sectors being non-metallic mineral products at 24.8% foreign content and refined petroleum at 49.9 per cent. Under the simple mean they count equally, giving 37.4%. Under export weighting, petroleum's much larger export value pulls the group figure toward its own share.


**Figure 5. Foreign value added share by skill group under alternative aggregation.** Panel A gives every sector equal weight when averaging sector-level shares. Panel B sums foreign value added and gross exports across sectors and countries before taking the ratio, so each sector carries weight proportional to its exports; vertical scales differ because pooled and equal-weighted levels are not comparable. Source: OECD TiVA database (2025 edition).

The panels differ in more than the ordering. In Panel A the two series are close to parallel: high-skill sits above low-skill in every year, and the gap holds at roughly 1.5 percentage points from 1995 through 2019. The ranking is a stable feature of the data, not an artefact of any particular year. In Panel B the ranking reverses and the gap is unstable — low-skill runs above high-skill from the start, widens to 3.6 points by 2012, then closes entirely, the two series meeting at 29.7 and 29.6 per cent in 2019.

That instability is itself informative. Export weighting hands most of the low-skill group to refined petroleum and basic metals, so the series inherits their commodity cycle: the sharp fall in 2009 and the decline after 2012 track the crude oil price rather than any change in how production is organised. The equal-weighted series carries no such exposure and rises steadily

throughout. The two panels are therefore not two views of one quantity. Panel A tracks how widely fragmentation has spread across sectors; Panel B tracks what the largest exporters do, and those exporters are disproportionately resource processors. Choosing between them is a choice about which question is being asked, which is why both are reported.

We can visualise this further by looking at the skill gap.

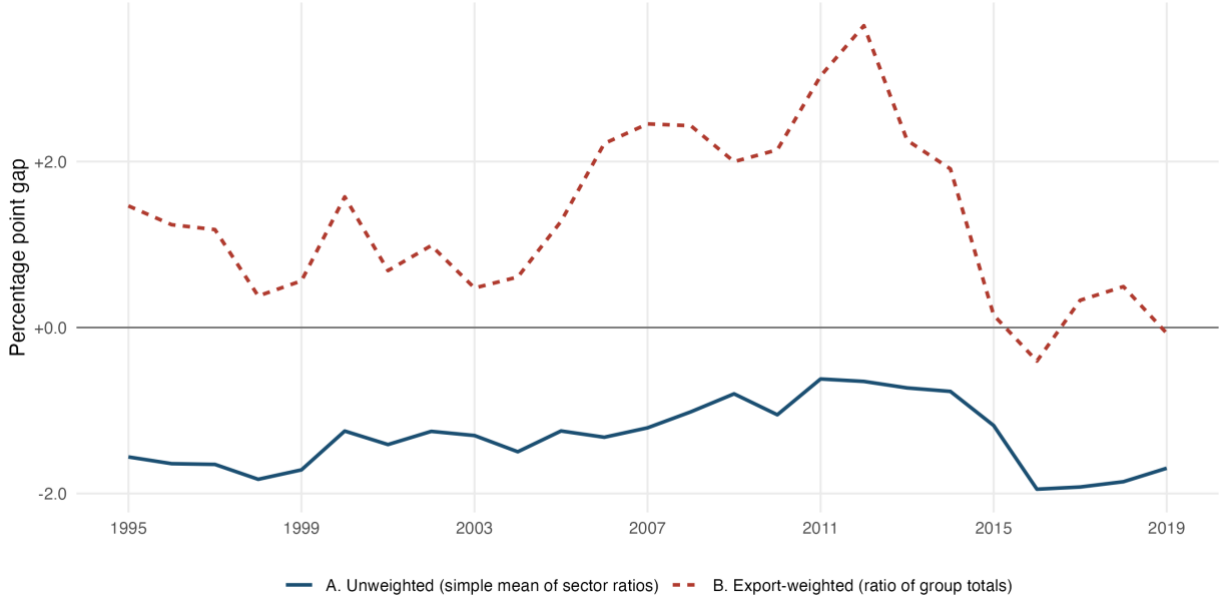


**Figure 6. The skill gap in foreign sourcing, low-skill minus high-skill.** The quantity plotted here is the regressor used in the headline specification. Source: OECD TiVA database (2025 edition).

The two series never cross. The unweighted gap is negative in every year, drifting between −1.6 and −0.7 before falling to −2.0 after 2014. The export-weighted gap is positive in 23 years of 25, climbing to +3.6 in 2012 and collapsing to zero by 2019. The two move together in shape, both rising to 2012 and falling after 2014, so they agree on when foreign sourcing shifted towards low-skill sectors. They disagree on whether it ever did. The export-weighted gap is the series with an aggregate interpretation, but it is dominated by refined petroleum and basic metals. Its post-2012 collapse tracks the crude oil price rather than any relocation of labour-intensive production. For a mechanism about displaced workers, that is the wrong exposure. The unweighted mean asks instead how far fragmentation has spread across sectors. Weighting by employment would settle it directly, but consistent sectoral employment data are unavailable across this panel.

A related caveat concerns the skill classification. Sectors are assigned to skill groups by technology intensity, following Hatzichronoglou (1997), which proxies for the skill composition of a sector's workforce rather than measuring it. The two can diverge: a capital-intensive sector may be classified as high technology while employing comparatively few skilled workers, and the taxonomy carries no wage information. Nothing in the classification identifies who is employed in a sector, only how research-intensive its output is. "Low-skill" should therefore be read throughout as "low technology intensity", and the displacement mechanism as operating through sectors whose production is technologically simple rather than through workers of an identified skill level. The employment-weighting extension proposed in Section 9 addresses both this and the weighting tension above.

## 2.6 Measuring populism

Conventional populism literature disagrees on how to define populism hence classifies it in different ways. The clearest definition is by Mudde (2004), society is divided into two homogeneous antagonistic groups, the pure people against the corrupt elite. Crucially populism is thin centred and attaches to a host ideology. Populism is traditionally measured using binary classification. Empirical work has typically operationalized this through binary classification, as in the PopuList (Rooduijn et al., 2019), Van Kessel, and the Chapel Hill Expert Survey. Each rest on expert coders judging which parties qualify, and judgement of this kind is not reducible to a rule (a second researcher applying the same criteria need not arrive at the same list). Guriev and Papaioannou (2022) caution that where classification depends on judgement, there is a standing temptation to apply the label to movements one happens to dislike. A score derived from a text-analysis rule applied uniformly to every manifesto avoids this. The rule itself remains open to challenge, but its application does not.

There are three further reasons to work with vote shares. The first is revealed preference. Voting is costly and it has consequences, so a vote tells us more than a survey answer does. Asking someone what they think about trade is cheap talk: nothing is at stake in the reply, and people tend to say what sounds acceptable. A vote is a choice made under real conditions. Secondly, populist parties in office change policy. Trade attitudes are a step removed from that. If the concern is what globalization does to politics, then the vote is the thing itself and not a stand-in for it. Third, it is continuous and it varies from one election to the next. A binary label would leave most of that variation out which is needed for a significant result.

Hence, I have chosen to measure on a continuous scale like Docquier et al. Docquier et al. (2026) examine the relationship between skill-specific globalisation shocks and populism. The key methodological contribution is a novel continuous populism score built from Manifesto Project party data. They capture two dimensions: the volume margin (share of votes for populist parties) and the mean margin (average populism score weighted by vote share). Low-skill imports increase right-wing populist vote shares while high-skill imports reduce it. Low-skill immigration increases right-wing populism but reduces left-wing populism.

Splitting the vote share by ideology is necessary to come up with any noteworthy mechanisms due to the nature of populism (not an ideology in itself) and due to the skill discrepancy, that exists within manufacturing. Right-wing populism runs on ethno-national grievance, left-wing on income and class. Since the mechanism under test is economic, the two should not respond identically, which is why they are estimated separately rather than pooled. Guriev and Papaioannou (2022) note the two differ substantively on redistribution.

When looking at the election data (number of elections) and populist vote share we find a right skewed distribution.

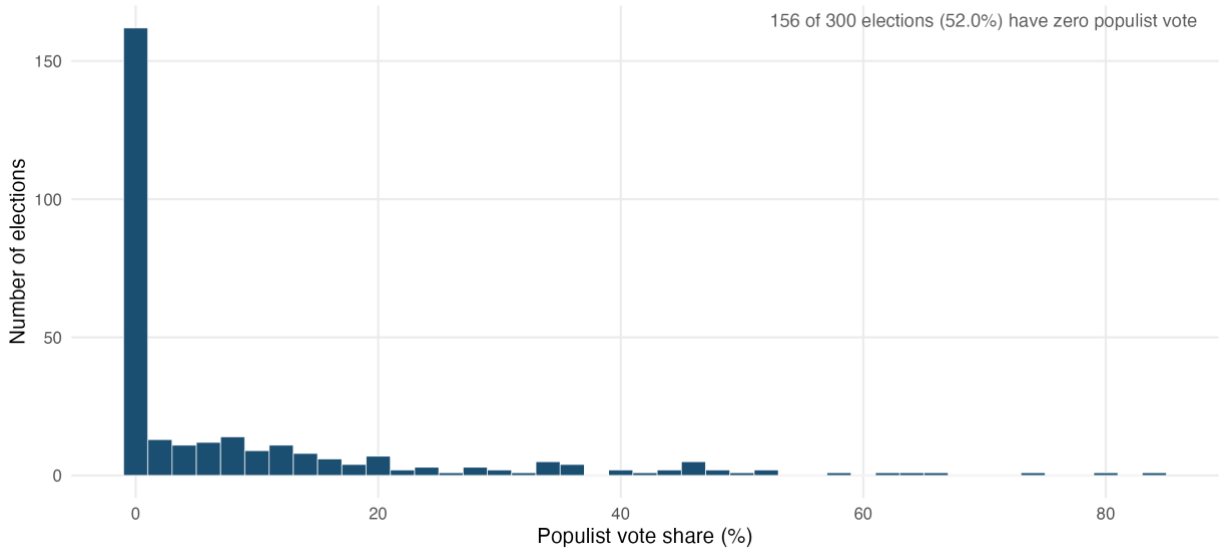


**Figure 7. Distribution of the populist volume margin.** 156 of 300 elections, or 52 per cent, record no populist vote. Source: author's calculations from party-level populism data.

The data shows that 156 of 300 elections, 52 per cent, record no populist vote at all. The remainder is heavily right skewed too. These zeros are informative rather than missing data as they are elections in which no party crossed the populism threshold. Hence, a log-linear specification would delete every one of them (ordinary least squares). This is the justification for PPML estimation method when using populist data.

Hence below the figure shows vote share by ideology.

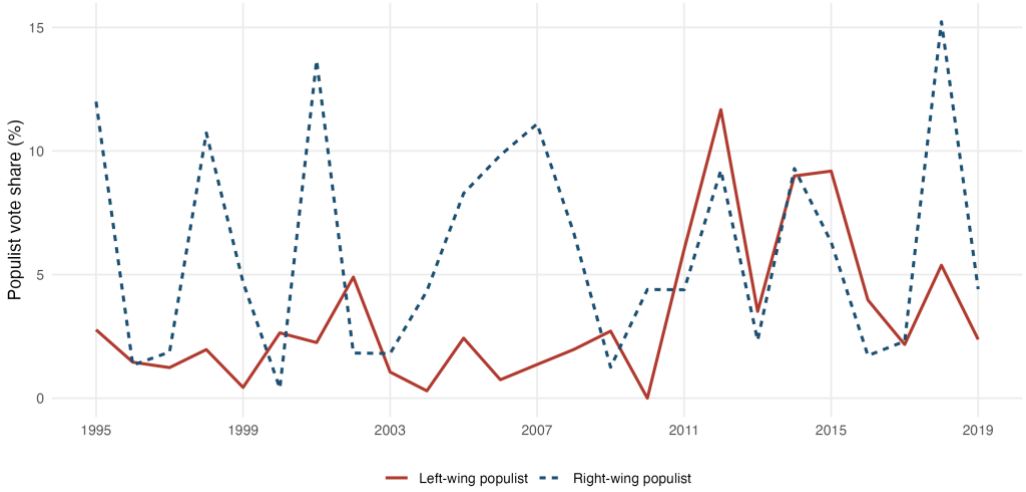


**Figure 8. Populist vote share over time, by ideological wing.** Parties are classified as right-wing where the rile score is positive and left-wing where it is zero or negative, following Laver and Budge (1992). Election timing is uneven across countries, so year-to-year movement partly reflects which countries held elections each year. Source: author's calculations from party-level populism data.

Right-wing populism dominates the late 1990s and 2000s, peaking near 13% in 2001 and 11% in 2007 while left-wing support sits near 2%. From roughly 2011 the two converge and trade places, with left-wing peaking at 11.7% in 2012. This is the same year the export-weighted skill gap peaks. Annual points average over whichever countries held elections that year, so individual spikes should not be over-read. In addition, when linking to TiVA must be careful as both show spikes in 2008, so rather than an exposure story could it just be a general economic grievance one from a financial shock.

The left-wing series also sits close to zero for the first half of the period, which means the left-wing estimates are identified largely from post-2011 variation. The export-weighted gap, by contrast, converges to essentially zero by 2019, so the 2012 co-movement is a coincidence of levels rather than a sustained relationship.

Taken together a mass of informative zeros, two wings that behave differently, and an exposure measure whose skill dimension is a technological proxy, these features of the data determine the specification set out below.

# 3. Methodology

This section sets out how exposure to foreign value added in exports is measured, and how that measure is related to populist voting. It begins with the construction of the exposure variable from OECD data and the aggregation of sector-level shares into the regressor used in estimation. It then defines the dependent variable and the estimating equation, before closing with the limitations of the approach.

## 3.1 Measuring exposure: domestic and foreign value added in exports

Following Koopman, Wang and Wei (2014), gross exports can be decomposed into their value-added origins using the inter-country input output framework. Production is interconnected: making one good requires inputs from other industries and other countries and those inputs required inputs of their own. Tracing value back to its origin therefore means accounting for the entire chain of upstream requirements not just the direct ones. The Leontief inverse does this, giving the total output required from every country and industry to produce one unit of final demand. Multiplying it by the value added coefficients converts that output into the value each country creates and the columns of the resulting matrix sum to one. Every unit of final output is fully attributed to the countries whose workers and capital produced the value within it. This is why the domestic and foreign shares defined below sum to exactly one hundred.

The OECD Trade in Value Added database publishes this decomposition directly. For exporting country c and industry i in year t, gross exports divide exhaustively into a domestic and a foreign component:

$$EXGR\_DVA_{c,i,t} + EXGR\_FVA_{c,i,t} = EXGR_{c,i,t} \quad (1)$$

Where EXGR denotes gross exports, EXGR_DVA the domestic value added embodied in them, and EXGR_FVA the foreign value added, all measured in current US dollars.

Domestic value added is the value created by factors located in country c, including value that left the country as an intermediate export, was processed abroad, and returned embedded in an import. Such re-imported value is correctly classified as domestic because it was created by domestic labour and capital and enters domestic GDP even if it has interacted within the valued chain through the various stages of production. Foreign value added is the residual value created outside country c and embodied in its exports.

Dividing through by gross exports gives the two share indicators:

$$DVASH_{c,i,t} = \frac{EXGR\_DVA_{c,i,t}}{EXGR_{c,i,t}} \times 100 \quad (2)$$

$$FVASH_{c,i,t} = \frac{EXGR\_FVA_{c,i,t}}{EXGR_{c,i,t}} \times 100 \quad (3)$$

which satisfy the exact identity

$$DVASH_{c,i,t} + FVASH_{c,i,t} = 100 \quad (4)$$

for every country, industry and year. FVASH is the measure of interest. It captures backward participation in global value chains: the proportion of a country's export value that originated

abroad. A manufacturing FVASH of forty means that forty cents in every euro of that country's manufacturing exports was created by foreign workers and firms upstream in the chain. This is precisely the quantity that gross trade statistics cannot recover, because they record the point at which goods cross borders rather than the point at which value is created.

Because DVASH and FVASH are linked by an exact identity, a regression on one is the linear mirror of a regression on the other. The coefficients would be complementary only differing in sign. FVASH is used throughout the analysis.

As we have seen in the motivation, FVASH is not a perfect variable and has to be analysed critically. As noted in Section 2, foreign value added share is high wherever purchased inputs are a large share of export value. This does not necessarily mean relocation of production.[7]

## 3.2 From sector shares to a country-level regressor

The OECD TiVA database reports FVASH for all industries. This has been restricted to the sixteen that fall under the manufacturing umbrella[8] (C in the ISIC Rev. 4 classification). Following Hatzichronoglou (1997), these sixteen are assigned to four tiers of technology intensity: high, medium-high, medium-low and low. Technology intensity is used here as a proxy for the skill of the workforce. The proxy is not perfect as skill intensity in this mechanism should reflect labour-intensive, low-wage employment, whereas technology intensity captures the sophistication of production, to follow the offshoring story explained in the motivation. A capital-intensive sector may be classified as high technology while employing few skilled workers, and the classification carries no information on wages. The four tiers are further collapsed into two skill groups, with high-tech and medium-high-tech forming the high-skill group and medium-low and low-tech forming the low-skill group[9].

Aggregating sector-level shares into a single figure for each group is a methodological choice rather than a mechanical step. Two rules are available, and in this panel, they lead to different interpretations so the choice must be made deliberately.

The first rule averages the sector ratios with equal weight, so that every industry counts once regardless of its size:

$$\text{FVASH}^{U}_{c,g,t} = \frac{1}{N_g} \sum_{i \in g} \text{FVASH}_{c,i,t} \tag{5}$$

where the superscript U denotes unweighted aggregation, g indexes the skill group, and N_g is the number of sectors in that group. Under this rule every sector, however small, has an equal say in the resulting FVASH figure.

[7] Referencing specifically coke and petroleum shown in the motivation.

[8] The sixteen codes used here are the finest non-overlapping partition available across the full country-year panel: C10T12, C13T15, C16T18, C19, C20, C21, C22, C23, C24, C25, C26, C27, C28, C29, C30 and C31T33. They differ in how narrowly they are drawn: some, such as C19 coke and refined petroleum, correspond to a single ISIC division, while others, such as C13T15 textiles, leather and footwear, aggregate three. Foreign content is therefore measured more precisely for some parts of manufacturing than others, and the broader categories may average over industries with quite different exposure. It was a methodological choice to include all available data to provide the best picture possible for an accurate representation of the manufacturing space.

[9] Collapsing to two groups is also warranted empirically, since the four tiers are highly correlated and no individual tier is significant when all four enter separately.

The second takes the ratio of group totals, which is how the published manufacturing aggregate is constructed:

$$\mathrm{FVASH}^{W}_{c,g,t} = \frac{\sum_{i \in g} \mathrm{EXGR\_FVA}_{c,i,t}}{\sum_{i \in g} \mathrm{EXGR}_{c,i,t}} \tag{6}$$

where W denotes export weighting, with each industry's share of group exports serving as its weight.

Substituting the definition of the sector share shows that the second rule is itself a weighted mean of the sector ratios, with each industry's share of group exports serving as its weight:

$$\mathrm{FVASH}^{W}_{c,g,t} = \sum_{i \in g} \omega_{c,i,t}\ \mathrm{FVASH}_{c,i,t},, \quad \omega_{c,i,t} = \frac{\mathrm{EXGR}_{c,i,t}}{\sum_{j \in g} \mathrm{EXGR}_{c,j,t}} \tag{7}$$

where ω is sector i's share of group exports.

A simple example shows why the choice matters. Suppose a group contains two sectors: non-metallic minerals, with 25% foreign content and exports of 10, and refined petroleum, with 50% foreign content and exports of 90. The simple mean gives 37.5%. The export-weighted mean gives 47.5% because the sector that exports most is also the sector with the highest foreign content. Had the shares been reversed, with petroleum at 25 and minerals at 50, export weighting would instead have returned 27.5%. The two rules therefore differ by the extent to which large exporters are also the foreign-intensive ones.

The choice of aggregation rule determines which skill group appears the more exposed. The equal-weighted mean is an average of ratios and corresponds to no observable aggregate. It is used nonetheless because the mechanism concerns how widely fragmentation has spread across sectors rather than how much export value it accounts for, and the export-weighted alternative is reported alongside it throughout.

The regressor used in the headline specification is the difference between the two skill groups:

$$\mathrm{GAP}_{c,t} = \mathrm{FVASH}_{c,\mathrm{low},t} - \mathrm{FVASH}_{c,\mathrm{high},t} \tag{8}$$

where low and high denote the low-skill and high-skill sector groups defined above.

A positive value indicates that low-skill manufacturing is more foreign-sourced than high-skill manufacturing, and an increase indicates that low-skill production has moved abroad faster than high-skill production. The two levels are strongly correlated, since both track a country's overall depth of value-chain participation, and entering them separately therefore identifies each coefficient from little residual variation. Differencing removes the common component and isolates the skill composition of foreign sourcing, which is the object the mechanism concerns.

## 3.3 The dependent variable

The dependent variable follows Docquier et al. (2026). For party p in country i contesting election e in year t, a populism score S is constructed by text analysis of party manifestos in the Manifesto Project Database. Two dimensions are extracted by principal components analysis: an anti-establishment stance and a commitment to economic protection. A party is classified as populist where its score exceeds a threshold τ, set at one standard deviation above the cross-party mean.

The volume margin is the combined vote share of all parties so classified:

$$\Pi^{V}_{i,e,t} = \sum_{p} \text{vote}^{p}_{i,e,t} \cdot \mathbb{1}\left[S^{p}_{i,e,t} > \tau\right] \quad (9)$$

where Π^V denotes the volume margin for country i at election e in year t, $vote^{p}$ is the vote share of party p in that election, $S^{p}$ is its populism score, τ is the classification threshold, and $\mathbb{1}[\cdot]$ is an indicator function taking the value one when the condition inside holds and zero otherwise.

It answers the question of how much of the electorate voted for a populist party, which is the quantity the displacement mechanism speaks to.

Docquier et al. also present a mean margin, defined as the vote-weighted average populism score across all parties. This measures the ideological centre of gravity of a party system. Since the mechanism under test concerns displaced workers casting votes, the volume margin is the more appropriate outcome and is used throughout. The mean margin captures something the volume margin does not, namely mainstream parties adopting populist positions without populist parties gaining seats and remains a natural extension.

Populism is thin-centred, attaching to a host ideology that determines which cleavage it mobilises. Since the mechanism under test is economic, left- and right-wing populism are not expected to respond identically to value-chain exposure. They are therefore estimated separately rather than pooled.

Parties are positioned on the right–left scale using the rile index in the Manifesto Project Database (Laver and Budge, 1992), following Docquier et al. (2026). A party is classified as right-wing where its rile score is positive and left-wing where it is zero or negative. Crossing this with the populism classification gives two wing-specific margins:

$$\Pi^{V,RW}_{i,e,t} = \sum_{p} \text{vote}^{p}_{i,e,t} \cdot \mathbb{1}\left[S^{p}_{i,e,t} > \tau\right] \cdot \mathbb{1}[\text{rile}^{p} > 0] \quad (10)$$

$$\Pi^{V,LW}_{i,e,t} = \sum_{p} \text{vote}^{p}_{i,e,t} \cdot \mathbb{1}\left[S^{p}_{i,e,t} > \tau\right] \cdot \mathbb{1}[\text{rile}^{p} \leq 0] \quad (11)$$

where $\text{rile}^{p}$ is party p's position on the Manifesto Project right–left scale, and the superscripts RW and LW denote the right-wing and left-wing margins respectively.

The split is exhaustive: every populist party is assigned to one wing or the other, so the two margins sum exactly to the pooled volume margin. Parties recording a rile score of precisely zero are assigned to the left. The estimation samples differ from the baseline, at 243 and 215 observations against 269, because elections in which no populist party of the relevant wing contested are absorbed by the fixed effects. The wing coefficients are consequently estimated on different sets of elections.

## 3.4 Empirical specification

The volume margin is zero in 125 of the 269 elections in the estimation sample, or 46 per cent, and the remaining mass is heavily right skewed. These features rule out the conventional

log-linear approach.[10] The volume margin is therefore estimated by Poisson pseudo-maximum likelihood:

$$\Pi^{V}_{i,e,t} = \exp\left[\alpha_i + \gamma_t + \beta_1 \cdot \mathrm{GAP}_{i,t} + \delta \cdot X_{i,e,t}\right] + \varepsilon_{i,e,t} \quad (12)$$

where $\alpha_i$ are country fixed effects, $\gamma_t$ are year fixed effects, GAP is the skill gap in foreign value added share defined above, X is a vector of controls and standard errors are clustered at country level. The exponential mean retains the zero observations as genuine data rather than deleting them.

Coefficients are semi-elasticities[11]. A one percentage point widening of the skill gap is associated with a $\beta_1 \times 100$ per cent change in the populist volume margin, expressed relative to the existing share rather than in percentage points.

The country fixed effects control for national characteristics that can affect populism, from political history to electoral structure and national identity. Whereas year fixed effects control for shocks that occurred in all countries simultaneously like a worldwide crisis (for example financial crisis 2008). What remains, and what identifies $\beta_1$, is variation within a country over time: whether a country records a higher populist vote in elections where its skill gap stands above its own historical average.

### *3.4.1 Controls*

Two controls are reported. The first is the log real exchange rate, proxied by the price level of consumption relative to the United States (Penn World Table 11.0). Real appreciation lowers the price of foreign inputs relative to domestic ones and therefore moves foreign value added share through a direct substitution channel. It also affects competitiveness and employment independently. The second is the log number of parties contesting the election following Docquier et al. (2026) because party fragmentation affects how vote share is distributed across a party system.

Several further controls were considered and rejected. Unemployment and the employment rate lie on the causal path running from offshoring through displacement to grievance, so conditioning on them would absorb part of the mechanism under test.

### *3.4.2 Limitations*

The mechanism under test concludes with displaced workers casting votes, but no employment data enters the analysis at any stage. What is observed is the composition of export value by skill intensity. Displacement is inferred from the established literature rather than measured here. The estimates should be read accordingly: they establish that a measure of exposure predicts voting, not that displacement has been observed to cause it.

---

[10] Taking logarithms of the dependent variable is undefined at zero, so every zero observation would be dropped. The retained subsample would not be random as it consists precisely of those elections in which populist parties were already present.

[11] The regressor enters in levels. Differencing was considered, since much of the related literature works in changes. Elections are spaced three to five years apart and unevenly across countries so a change measured since the previous election would span different intervals across observations. The country fixed effects extract within-country variation without imposing a fixed interval.

FVASH can only be computed for sectors that continue to export. Industries displaced outright by foreign competition leave the data altogether rather than registering as high-exposure observations. The measure structurally omits the most severely affected cases. The resulting bias understates exposure and pushes the coefficient towards zero. A significant estimate therefore survives a design working against it and should be read as a conservative lower bound. The alternative case, in which the bias inflated the estimate, would be considerably more damaging and does not arise here.

The estimates are conditional correlations rather than causal effects. Reverse causality cannot be ruled out: populist parties in office may enact policies that themselves reduce value-chain participation. A shift-share instrument constructed from origin-country supply shocks would address this but requires the additional assumption that such shocks affect value-chain participation through the same channel as gross trade, which is plausible but less firmly established than in the import-competition setting.

# 4. Results

The results begin with the simplest specification.

## Table 1. Foreign value added and the populist volume margin, PPML

| | (1) Skill split | (2) Gap | (3) Right-wing | (4) Left-wing |
|---|---|---|---|---|
| FVASH low skill | 0.111** | | | |
| | (0.050) | | | |
| FVASH high skill | −0.094* | | | |
| | (0.057) | | | |
| FVASH gap (low minus high) | | 0.101** | 0.059 | 0.303*** |
| | | (0.049) | (0.088) | (0.088) |
| Observations | 269 | 269 | 243 | 215 |
| Countries | 39 | 39 | 35 | 32 |
| Pseudo $R^2$ | 0.416 | 0.416 | 0.370 | 0.543 |
| Country FE | ✓ | ✓ | ✓ | ✓ |
| Year FE | ✓ | ✓ | ✓ | ✓ |

Notes: ***, **, and * denote significance at the 1, 5 and 10 per cent levels respectively; clustered standard errors at the country level are reported in parentheses; all coefficients estimated with PPML using the R command fepois. The dependent variable in columns (1) and (2) is the populist volume margin; in columns (3) and (4) it is the volume margin restricted to parties with positive and non-positive rile scores respectively. Columns (3) and (4) rest on fewer elections because those in which no populist party of the relevant wing contested are absorbed by the country fixed effects. All coefficients are reported on the FVASH scale.

The skill-split specification gives 0.111 for low-skill manufacturing, significant at the 5 per cent level, and −0.094 for high-skill. The two enter with opposite signs. PPML coefficients are semi-elasticities, so a one percentage point rise in low-skill foreign content is associated with an increase of 11.7% in the populist volume margin. That is 11.7% of the existing vote share rather than eleven percentage points.[12]

Over the period a rising share of the value in low-skill manufacturing exports was created abroad. This is offshoring: stages of production once performed domestically are now performed elsewhere and imported back as inputs. The measure captures the shift in where production takes place, not the employment consequence directly, and the inference drawn from the literature is that the stages leaving high-wage economies are the labour-intensive ones. The losses are concentrated among the workers who performed them, while the gains are diffuse and accrue elsewhere.

Populism itself is not a programme. Being thin-centred, what it offers is not a policy but a diagnosis: a division between an ordinary people and an elite held responsible for their circumstances. That diagnosis fits the experience of workers whose sector was relocated by decisions taken above them and defended by parties of government. A vote for such a party is closer to a revealed preference over grievances than a choice between economic platforms, which is why exposure predicts the vote without any policy specifically being on offer[13].

[12] a country polling 10% for populist parties would move to about 11.1%

[13] These figures use the simple mean, which weights every sector equally regardless of size. The measure is sensitive to small surviving industries within a country and low-skill and high-skill FVASH move colinearly

Due to the issues with the simple model, the gap specification is preferred, but this does not take away from the possible economic explanation. The gap specification combines the two variables into one. The gap coefficient is 0.101, again significant at the 5 per cent level, so a one percentage point widening of the gap is associated with an increase of about 10.6% in the populist volume margin. But once we introduce the gap, a one point widening could mean low-skill rose, high-skill fell, or both moved. Those are different economic events with the same coefficient. A reader can reasonably ask which is driving it.

The volume margin can be decomposed by wing using the rile index, as set out in the methodology. One caveat applies before the results are read. The wing samples contain 243 and 215 observations against a baseline of 269, since elections in which no populist party of the relevant wing contested are absorbed by the fixed effects. The three columns therefore rest on different sets of elections.

Both wing specifications use the skill gap as the regressor, and the effect is concentrated entirely on the left. The left-wing coefficient is 0.303 and significant at the 1% level. The right-wing coefficient is 0.059 and indistinguishable from zero. The left-wing coefficient of 0.303 implies that a one percentage point widening of the gap is associated with an increase of about 35% in the left-wing populist vote share.

The offshoring account explains the left-wing result directly. Displaced low-skill manufacturing workers sit in the lower half of the income distribution, so the grievance offshoring generates is material in origin. Populism itself supplies no policy but is adjunctive to the party's original ideology. On the left that host is redistributive, offering higher taxation, stronger social protection and hostility towards firms that relocate production. Rodrik (2021) frames the pattern as a distinction between cleavages: shocks arriving through trade and investment activate the income cleavage and mobilise the left, whereas shocks arriving through immigration activate the ethno-national cleavage and mobilise the right.

The same mechanism cannot explain the right-wing null. The null is consistent with right-wing populism mobilising a different cleavage. Its voters may respond to immigration rather than to the relocation of industry, in which case a measure built from foreign value added cannot detect the shock they are reacting to. The framing of a national people against outsiders is the lever available to right-wing populists, and it is not one that offshoring supplies. A second possibility is that its support is drawn from constituencies offshoring leaves untouched. Losses through manufacturing are concentrated on certain regions but this cannot explain the full populism trend. Populism is a global trend so other areas not suffering from job losses through trade will vote populist too but will not be picked up by this mechanism.

A further concern is heterogeneity across political systems. The panel spans countries with very different party structures, electoral rules and populist traditions, and the fixed effects absorb differences in level rather than in responsiveness. Argentina, for instance, has a long Peronist tradition in which populist appeals are a standing feature of political competition rather than a response to economic shocks. If the relationship between exposure and voting is stronger in some systems than others, the estimate is an average that may not describe any of them well.

A related limitation concerns mainstream parties. Where populist sentiment grows, established parties may adopt populist positions rather than cede votes to new entrants. The volume margin counts only parties crossing the populism threshold, so accommodation of this

kind is invisible to it. The mean margin would capture it, which is one reason to treat that measure as a natural extension.

The specification includes country fixed effects, and Figure 9 shows why they are necessary rather than conventional.

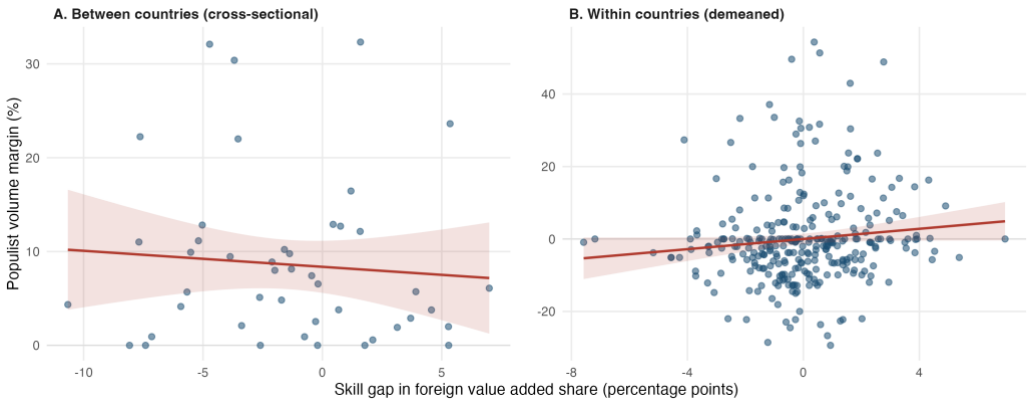


**Figure 9. The skill gap and populist voting, between versus within countries.** The cross-sectional slope is −0.17 and the within-country slope is +0.70[14]. Source: OECD TiVA (2025 edition) and author's calculations.

The two panels ask different questions of the same data. Panel A compares countries with one another: each point is a country, plotting its average skill gap against its average populist vote share. On this comparison, countries with wider gaps are if anything less populist, with a fitted slope of −0.17. Panel B compares each country with itself, plotting every election as a deviation from that country's own average, so a point to the right of zero marks an election in which the gap was wider than usual and a point above zero one in which the populist vote was higher than usual. Here the relationship reverses, sloping upward at +0.70.

The reason for the reversal is that between-country differences in foreign sourcing mostly reflect country size. Small open economies import inputs by necessity, so they record high foreign value added shares regardless of any offshoring decision. In this panel they also happen to record lower populist vote shares. The cross-sectional association therefore captures a compositional fact about the sample rather than a behavioural relationship. Country fixed effects remove it entirely, leaving only the comparison of a country with itself over time. This is the clearest justification available for the design. Without fixed effects, the same data would have produced completely different results.

As explained in the motivation and method, it is important to not dig too deeply into the simple means due to the bias created although the possible economic explanations remain. Therefore, I constructed an export weighted variable.

[14] Both fitted lines are ordinary least squares on the raw data and are shown to indicate direction they are not comparable with the Poisson estimates reported above. The panels also include all 300 observations, whereas the estimates use 269, since six countries recording no populist vote in any election are absorbed by the fixed effects.

## Table 2. Sensitivity to the aggregation scheme

| Specification | Unweighted | Export-weighted |
|---|---|---|
| Skill split (low skill) | 0.111** | 0.053 |
| | (0.050) | (0.036) |
| **Gap (low minus high)** | **0.101**** | **0.057*** |
| | (0.049) | (0.033) |
| Observations | 269 | 269 |

Notes: PPML with country and year fixed effects, standard errors clustered at country level. All coefficients on the FVASH scale.

Under export weighting a one percentage point widening of the skill gap is associated with an increase of 5.9% in the predicted populist vote, against 10.6% under the simple mean. Both remain significant, at 10% and 5% respectively. The skill-split coefficients do not survive the change of weighting, but the gap does.

The economic content of the measure changes under export weighting. A handful of very large sectors dominate the weighted aggregate, and those sectors are typically capital-intensive or resource-processing: refined petroleum above all, alongside basic metals and chemicals. Their foreign content arises from imported raw materials and from the internal sourcing of multinational firms rather than from production stages relocating to lower-wage economies. Therefore, the weighted measure captures a different phenomenon, one closer to resource dependence and transnational production structure than to offshoring in the sense the mechanism requires. The simple mean has the opposite property. Every industry counts once, so exposure rises when foreign sourcing spreads across many low-skill sectors rather than when one large exporter shifts to imported inputs.

The fall in the coefficient is therefore unsurprising. Export weighting moves the measure away from a labour market grievance story, so a weaker estimate is what one would expect. That it remains positive and significant means the headline result does not rest on the contested choice. These estimates should not be read as a full account of how global value chains relate to populism. Exports are a single margin within the value-added space, and the relationship identified here is a specific one, between the skill composition of foreign sourcing and populist voting.

Including the controls, we get similar results.

## Table 3. Controls and robustness

| | (1) Baseline | (2) + Exch. rate | (3) + Parties | (4) Both | (5) No crisis |
|---|---|---|---|---|---|
| FVASH gap (low minus high) | 0.101** | 0.101** | 0.100** | 0.100** | 0.070 |
| | (0.049) | (0.049) | (0.049) | (0.049) | (0.055) |
| Log real exchange rate | | −0.056 | | −0.023 | −0.328 |
| | | (0.888) | | (0.879) | (1.081) |
| Log number of parties | | | 0.084 | 0.081 | 0.190 |
| | | | (0.451) | (0.450) | (0.563) |
| Observations | 269 | 269 | 269 | 269 | 206 |
| Pseudo $R^2$ | 0.416 | 0.416 | 0.416 | 0.416 | 0.388 |

| | | | | | |
|---|---|---|---|---|---|
| Country FE | ✓ | ✓ | ✓ | ✓ | ✓ |
| Year FE | ✓ | ✓ | ✓ | ✓ | ✓ |

Notes: ***, **, and * denote significance at the 1, 5 and 10 per cent levels respectively; clustered standard errors at the country level are reported in parentheses; all coefficients estimated with PPML using the R command fepois. Column (5) excludes elections held between 2008 and 2012. All coefficients are reported on the FVASH scale.

The estimate is invariant across specifications, moving from 0.101 to 0.100 with the observation count unchanged at 269 in columns (1) to (4). Both controls are individually insignificant.

Each addresses a specific alternative explanation. The real exchange rate matters because appreciation lowers the price of foreign inputs relative to domestic ones, which mechanically raises foreign value added share without any decision to relocate production. It also affects competitiveness and employment directly so it could plausibly drive both the regressor and the outcome. The number of parties matters because vote shares are distributed across whatever parties contest an election: a populist party in a three-party system holds a larger share than the same party in a ten-party system independently of how much support it commands. Neither control moves the estimate so neither explains the relationship.

Column (5) excludes elections held between 2008 and 2012, addressing the possibility that the financial crisis rather than offshoring drives the result. The concern is well founded. Funke et al. (2016) document that financial crises are systematically followed by rising support for far-right parties, and the bailouts of 2008 provided populist parties with a direct target: public money committed to institutions widely held responsible for the crisis. Since value chain integration deepened over the same period, the two channels are difficult to separate. The mechanism at play can be explained trivially. When people lose purchasing power or suffer from a crisis, they start to experience economic grievances. These grievances can be reflected in revealed preferences in elections by voting against the status quo or for populist parties. Populist messaging is effective at capturing this electorate due to the antiestablishment stance.

In addition, how countries react to a crisis can explain this too economically. If there is a global financial crisis the year fixed effects should absorb them however it does not absorb the interactions between countries. A financial crisis would propagate throughout the valued chain as countries would strip back on exports and focus on domestic production. When exports fall, the openness of the valued chain reduces, so the foreign content of exports share should fall too as those exports would be replaced with local production. The relationship between FVASH and economic crises therefore compounds over time.

The coefficient falls to 0.070 and loses conventional significance, but 63 observations are dropped in the process. The point estimate remains positive and around 70 per cent of the baseline. The export-weighted skill gap and left-wing populist support both peak in 2012, so removing those years removes a substantial part of the identifying variation. Nevertheless, the crisis channel matters for populist support, and conclusions should be drawn with that in mind.

# 5. Discussion and Conclusion

To conclude, when interpreting foreign value added in exports and its relationship with populism, the specification is key for possible economic interpretations. The simple mean FVASH can sell an offshoring story, whereas when we compute the export weighted mean it is driven by natural resource sectors within manufacturing. In addition, interpreting the skill gap is more difficult than the simple skill split due to the unknown of whether it is high skill or low skill driving the higher or lower skill gap. Splitting foreign value added into its high-skill and low-skill components is conceptually the sharper test, but it is the least well identified. The two series are not independent: both the high- and low-skill components of FVASH trend upwards across the sample so the two regressors share most of their variation. What remains once the common movement is absorbed is a small residual, and the standard errors widen accordingly. The consequence is a loss of statistical power rather than a change in the point estimates and the intervals are wide enough that the two cannot be cleanly separated from each other.

This is why the skill gap is the better-behaved specification even though it is the harder one to interpret. Differencing removes the common trend the two levels share, so the remaining variation is the composition of foreign content rather than its level, and the estimate is recovered from variation that the two-way split spends on collinearity.

A second limitation concerns which sectors appear in the data at all. FVASH is a ratio of foreign value added to gross exports, so it exists only for country-sector pairs that still export. An industry that fragmented until domestic production ceased does not enter the panel as an extreme observation of exposure, it leaves the panel. The sectors available to be measured are, by construction, those that adapted to fragmentation rather than those destroyed by it.

The direction of the resulting bias is the point. The missing observations are precisely the highest-exposure ones, and their absence compresses the upper tail of the explanatory variable. Since the political response should be strongest exactly where displacement was most complete, the coefficient is pushed towards zero. The estimates here are therefore a conservative lower bound on the relationship between fragmentation and left-wing populist support, and a significant coefficient has survived a design working against it.

However, it is clear to see even if it more of a correlation than causation, a relationship between foreign value added and left wing populism exists. As explained above, offshoring using the simple mean, across sectors, is more concentrated in low skill sectors. These sectors tend to be more labour intensive. Hence, the losses are concentrated on workers in certain key regions. These workers express their economic grievances[15] in elections through revealed preferences by voting against the status quo. Populist rhetoric expresses an antiestablishment stance so can easily attract these workers. This works in conjunction with the fact these workers are more likely to be working class so are also attracted by left sided policies too. Hence, offshoring is focused on labour intensive sectors and can increase populist support in left winged parties. As established before, the model cannot explain the right winged populist increase in support meaning it is likely from a different mechanism.

This evidence is supported by Rodrik's findings. Rodrik argues globalisation shocks activate one of two political cleavages. The income cleavage creates winners against losers in economic terms and is activated by trade, finance and foreign investment shocks. This

[15] For example, loss of employment with strong labour market frictions

mobilises the left because the natural response is redistribution and protection. The ethnical/national cleavage being natives against outsiders, is activated by immigration and refugee shocks. This mobilises the right because the natural response is exclusion. FVASH is a production fragmentation measure. Therefore, it is a trade and investment shock, so Rodrik predicts it activates the income cleavage.

The apparent tension with Colantone and Stanig (2018) is instructive rather than damaging because the two papers do not measure the same shock. They measure import competition, an output-market shock to firms competing with those imports. This paper measures the foreign content of a country's own exports, which is an input-sourcing shock to firms participating in the same production network. In the first case the domestic firm loses the sale and in the second it keeps the sale but relocates part of the process that produced it. These are different events for a worker, and following Rodrik's (2018) logic, they should activate different cleavages. Import competition arrives as goods from a foreign country and is readily narrated as a national loss to a named rival. That framing mobilises the ethnic-national cleavage and the parties positioned on it. Offshoring arrives as a decision by a domestic employer about where to place a stage of production: there is no foreign actor to blame and there is a visible distributional counterpart at home. The divergence is a contribution to how the literature partitions trade shocks rather than a contradiction of it. A rightward result from import competition and a leftward result from fragmentation are jointly consistent if the mechanism is cleavage activation and the two measures activate different cleavages.

Autor et al. (2020) support that reading from a third direction. They find that trade exposure polarises rather than uniformly shifts. Exposed districts elect more extreme representatives in both directions, with the direction determined by local demographic composition rather than by the shock itself. A single backlash mechanism cannot generate that pattern. A framework in which exposure raises the salience of a cleavage, and pre-existing local conditions determine which cleavage is available to be mobilised, can.

Docquier et al. (2026) are the closest available comparison. They find that low-skill labour-intensive imports raise right-wing populism while leaving left-wing populism largely unchanged and conclude that globalisation cannot be blamed. What matters is the skill content of the shock. They measure trade in gross terms. The export side of the skill-content story activates the income cleavage while the import side activates the national one. The left-wing margin is thus responsive to trade exposure in their data too, its sign depending on which exposure is measured and on what institutions are available to absorb it.

We find that the financial crisis complicates the interpretation. However, deep GVC integration means a country imports other countries' crisis too. A shock anywhere in the chain propagates to everyone downstream, so an integrated economy faces not only its own shock but its trade partners' as well. The two channels are not competing explanations but compounding ones. Bems, Johnson and Yi (2013) document that GVC trade amplified the 2009 collapse supporting this explanation.

The most serious threat to interpretation is reverse causality, and it runs through policy. The regressions relate foreign value added in exports to populist vote shares, but populist parties in office have the instruments to reduce foreign value added. Populists in office does not guarantee anti-globalisation policy, but it makes it likely. One of the two dimensions on which the populism score is built upon is a commitment to protect the nation's ordinary people. An easy way to make that commitment is by restricting cross-border production. This can come in the form of tariffs on intermediate inputs, local content requirements, procurement rules favouring domestic suppliers, restrictions on inward and outward foreign

direct investment and so on and so forth. The parties whose vote share is the dependent variable are precisely those most likely to act on the explanatory variable.

What is not available is a control for policy stance itself. No panel measure of anti-globalisation policy covers this sample at the required frequency. The honest statement is that these are conditional correlations rather than causal effects and that the shift-share design proposed below is the natural route to separating them.

Measuring exposure on the export side was defended on three grounds, and the costs are best set out alongside them. The justification was, first, that the accounting supports it, the foreign value added decomposition in TiVA is defined on gross exports, so the export bundle is the only place a clean share of foreign content can be recovered rather than approximated. Second, that fragmentation is conceptually an export-side phenomenon: offshoring is observed when a country's own sales embody stages performed elsewhere, which is what the foreign share of exports records. Import volumes record competition without revealing where a domestic producer's own process is located. Third, that exporting sectors are the exposed tradable sectors and those in which relocation decisions are taken. Wage and employment consequences concentrate and are geographically clustered enough for local labour-market effects to register electorally.

Two things this misses. Offshoring by firms serving only the domestic market is invisible: a producer that relocates a stage of production but sells entirely at home displaces the same workers and contributes nothing to FVASH. The same fragmentation logic applies to non-exporters, so the exposure variable understates total offshoring everywhere. It understates it most where domestic-market production is large relative to trade. Second, the measure is less informative for economies with small export sectors, where a handful of large exporters can dominate the aggregate and the resulting share describes those firms rather than the national production structure. Both omissions run in the same direction as the survivorship problem. The exposure variable is systematically undervalues the exposure workers face, and in all three cases the measurement error takes the form of missing high-exposure mass rather than added noise around a correct value.

On average, only around 60% of what a country exports is its own. The remaining share is value created elsewhere and carried across a border on its way to a final market. That single fact is why gross trade statistics are the wrong instrument for measuring how globalisation is experienced: a country's export figures describe what it sells, not what it makes, and the gap between the two is where fragmentation lives.

This paper's contribution follows from taking that gap seriously in a political setting. The existing literature on globalisation and populism measures exposure in gross terms. Measuring the skill content of exports in value-added terms changes what is being asked. A gross export measure attributes an entire shipment to the exporting country regardless of how much of it that country produced. A value-added measure separates domestic from foreign content and registers the international division of labour within a product. Adding the skill dimension to that decomposition turns the question from how much a country trades into which kinds of work it has stopped doing at home. Docquier et al. (2026) come closest in taking the skill content of globalisation seriously, but they measure it on the import side, in gross terms, and set exports aside explicitly.

It is clear to see how you measure globalisation exposure determines what you conclude about who it hurts. The same countries, the same years and the same populism data yield a rightward result when exposure is import competition and a leftward one when exposure is

the foreign content of exports. Additionally, offshoring and import competition activate different politics. Fragmentation, concentrated as it is in low-skill labour-intensive sectors, produces grievances with a domestic distributional shape and a redistributive answer, which the populist left is positioned to supply. Import competition arrives as a foreign country's goods and invites a national answer, which the populist right supplies. Treating the two as interchangeable proxies for the same thing obscures a distinction voters appear to be making.

Several extensions follow from the limitations above. The most direct concerns aggregation: sector shares are weighted here by export value or with the unweighted mean. A mechanism operating through labour displacement would ideally weight by employment. Consistent sectoral employment data are unavailable across this panel, but a narrower sample of European countries would permit both employment weighting and a direct test of the displacement step that this paper can only infer.

A second extension concerns identification. A shift-share instrument for foreign value added, constructed from origin-country supply shocks, would address the possibility that populist governments themselves reduce value-chain participation, though it requires the assumption that such shocks affect fragmentation through the same channel as gross trade.

Finally, the political side merits extension. The right-wing result here is imprecisely estimated and consistent with zero, which reflects my mechanism not being able to explain its variation. It would be interesting to compute a working specification to capture right wing populism.

# 6. Data

The analysis combines two sources in a country-election panel. The political data come from Docquier et al. (2026), who construct populism measures from the Manifesto Project Database for 55 countries between 1960 and 2019, covering 4,140 party-election observations across 1,132 parties. Each party contesting an election receives a populism score derived from text analysis of its manifesto, combining an anti-establishment index with an index of commitment to economic protection. A party is classified as populist where its score exceeds the cross-party mean by more than one standard deviation, and the vote shares of all such parties are summed to form the outcome variable. Parties are separately positioned on the rile left–right index, which permits the wing-specific margins defined in Section 3.

Exposure is measured using the 2025 edition of the OECD Trade in Value Added database, which reports the foreign value added share of gross exports at country-industry level for 45 countries between 1995 and 2019. The sixteen manufacturing divisions of ISIC Rev. 4 are assigned to four technology-intensity tiers following Hatzichronoglou (1997), then collapsed into two skill groups: high-skill comprises the high and medium-high tiers, low-skill the medium-low and low tiers. Aggregating within each group yields the two regressors from which the skill gap is constructed.

The panel begins in 1995, the first year of TiVA coverage, and ends in 2019, before the pandemic disrupted value chains for reasons unrelated to the mechanism studied. Merging on country and election year gives 300 election observations with complete exposure data. Estimation drops a further six countries and 31 observations, in every case because the country records no populist vote in any election in the window. The estimation sample is therefore 269 observations across 39 countries.

The six are Estonia, Luxembourg, Malta, Norway, Sweden and the United States. Their exclusion follows mechanically from Poisson estimation with country fixed effects: where the outcome is zero throughout. The fixed effect fits it perfectly and the observations carry no identifying variation. It also reflects the strictness of the underlying classification which flags a party as populist only where its manifesto score exceeds the cross-party mean by more than one standard deviation.

## Table A.1. Manifesto Project Database — estimation sample

| | i | 1st E. | Last E. | # E. | # P. | | i | 1st E. | Last E. | # E. | # P. |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Denmark | 1 | 1960 | 2019 | 22 | 21 | Greece | 24 | 1974 | 2019 | 18 | 21 |
| Japan | 2 | 1960 | 2017 | 20 | 30 | Portugal | 25 | 1975 | 2019 | 16 | 22 |
| New Zealand | 3 | 1960 | 2017 | 20 | 14 | Spain | 26 | 1977 | 2019 | 15 | 46 |
| Sweden | 4 | 1960 | 2018 | 18 | 16 | Argentina | 27 | 1989 | 2013 | 6 | 14 |
| USA | 5 | 1960 | 2016 | 15 | 2 | Chile | 28 | 1989 | 2017 | 6 | 15 |
| Australia | 6 | 1961 | 2019 | 23 | 14 | Bulgaria | 29 | 1990 | 2017 | 10 | 28 |
| Belgium | 7 | 1961 | 2019 | 18 | 36 | Croatia | 30 | 1990 | 2016 | 9 | 35 |
| Germany | 8 | 1961 | 2017 | 16 | 11 | Czech Rep. | 31 | 1990 | 2017 | 9 | 25 |
| Ireland | 9 | 1961 | 2016 | 16 | 16 | Hungary | 32 | 1990 | 2018 | 8 | 17 |
| Israel | 10 | 1961 | 2019 | 18 | 64 | Romania | 33 | 1990 | 2016 | 8 | 31 |
| Mexico | 11 | 1961 | 2018 | 20 | 26 | Slovakia | 34 | 1990 | 2016 | 9 | 27 |
| Norway | 12 | 1961 | 2017 | 15 | 13 | Slovenia | 35 | 1990 | 2018 | 9 | 26 |
| Turkey | 13 | 1961 | 2018 | 16 | 22 | Poland | 36 | 1991 | 2019 | 9 | 37 |
| Austria | 14 | 1962 | 2019 | 18 | 13 | Estonia | 37 | 1992 | 2019 | 8 | 26 |
| Canada | 15 | 1962 | 2019 | 19 | 9 | Lithuania | 38 | 1992 | 2016 | 7 | 27 |
| Finland | 16 | 1962 | 2019 | 16 | 20 | Sth Korea | 39 | 1992 | 2016 | 7 | 16 |
| France | 17 | 1962 | 2017 | 14 | 35 | Latvia | 40 | 1993 | 2018 | 9 | 36 |
| Iceland | 18 | 1963 | 2017 | 17 | 19 | Russia | 41 | 1993 | 2011 | 6 | 25 |
| Italy | 19 | 1963 | 2018 | 15 | 58 | Sth Africa | 42 | 1994 | 2019 | 6 | 13 |
| Netherlands | 20 | 1963 | 2017 | 17 | 28 | Ukraine | 43 | 1994 | 2019 | 8 | 41 |
| Switzerland | 21 | 1963 | 2019 | 15 | 25 | Cyprus | 44 | 1996 | 2016 | 5 | 12 |
| Luxembourg | 22 | 1964 | 2013 | 11 | 11 | Malta | 45 | 1996 | 1998 | 2 | 2 |
| UK | 23 | 1964 | 2019 | 16 | 14 | | | | | | |
| | | | | | | **Total** | | | | **585** | **1059** |

Notes: Adapted from Docquier et al. (2026), Table A.1. Columns report the first and last election observed, the number of elections, and the number of distinct parties, with totals recomputed over the retained sample. The 45 countries listed are those common to the Manifesto Project sample and the OECD TiVA database; ten of the original 55 are excluded for lack of TiVA coverage: Albania, Armenia, Azerbaijan, Bolivia, Bosnia-Herzegovina, Georgia, Moldova, Montenegro, North Macedonia and Serbia. A further six countries are absorbed in estimation because they record no populist vote in any election, leaving 269 observations across 39 countries. No observations are lost as singletons.

# 7. References


Antràs, P., de Gortari, A. and Itskhoki, O. (2017). Globalization, inequality and welfare. Journal of International Economics, 108, 387–412. https://doi.org/10.1016/j.jinteco.2017.07.009

Autor, D., Dorn, D., Hanson, G. and Majlesi, K. (2020). Importing political polarization? The electoral consequences of rising trade exposure. American Economic Review, 110(10), 3139–3183. https://doi.org/10.1257/aer.20170011

Bems, R., Johnson, R. C. and Yi, K.-M. (2013). The great trade collapse. Annual Review of Economics, 5, 375–400. https://doi.org/10.1146/annurev-economics-082912-110201

Colantone, I. and Stanig, P. (2018). The trade origins of economic nationalism: import competition and voting behavior in Western Europe. American Journal of Political Science, 62(4), 936–953. https://doi.org/10.1111/ajps.12358

Docquier, F., Iandolo, S., Rapoport, H., Turati, R. and Vannoorenberghe, G. (2026). Populism and the skill-content of globalization. The Economic Journal, advance access, published online 16 February 2026. https://doi.org/10.1093/ej/ueag019

Funke, M., Schularick, M. and Trebesch, C. (2016). Going to extremes: politics after financial crises, 1870–2014. European Economic Review, 88, 227–260. https://doi.org/10.1016/j.euroecorev.2016.03.006

Guriev, S. and Papaioannou, E. (2022). The political economy of populism. Journal of Economic Literature, 60(3), 753–832. https://doi.org/10.1257/jel.20201595

Hatzichronoglou, T. (1997). Revision of the high-technology sector and product classification. OECD Science, Technology and Industry Working Papers, No. 1997/02 (OCDE/GD(97)216). OECD Publishing, Paris. https://doi.org/10.1787/134337307632

Helpman, E. (1984). A simple theory of international trade with multinational corporations. Journal of Political Economy, 92(3), 451–471. https://doi.org/10.1086/261236

Johnson, R. C. (2018). Measuring global value chains. Annual Review of Economics, 10, 207–236. https://doi.org/10.1146/annurev-economics-080217-053600

Johnson, R. C. and Noguera, G. (2012). Accounting for intermediates: production sharing and trade in value added. Journal of International Economics, 86(2), 224–236. https://doi.org/10.1016/j.jinteco.2011.10.003

Koopman, R., Wang, Z. and Wei, S.-J. (2014). Tracing value-added and double counting in gross exports. American Economic Review, 104(2), 459–494. https://doi.org/10.1257/aer.104.2.459

Laver, M. and Budge, I. (eds.) (1992). Party Policy and Government Coalitions. Macmillan, London. https://doi.org/10.1007/978-1-349-22368-8

Mudde, C. (2004). The populist zeitgeist. Government and Opposition, 39(4), 541–563. https://doi.org/10.1111/j.1477-7053.2004.00135.x

OECD (2025). Trade in Value Added (TiVA) database, 2025 edition. OECD, Paris. Available at https://www.oecd.org/en/topics/sub-issues/trade-in-value-added.html (accessed 17 August 2026).

Rodrik, D. (2018). Populism and the economics of globalization. Journal of International Business Policy, 1(1–2), 12–33. https://doi.org/10.1057/s42214-018-0001-4

Rodrik, D. (2021). Why does globalization fuel populism? Economics, culture, and the rise of right-wing populism. Annual Review of Economics, 13, 133–170. https://doi.org/10.1146/annurev-economics-070220-032416

Rooduijn, M., Pirro, A. L. P., Halikiopoulou, D., Froio, C., Van Kessel, S., De Lange, S. L., Mudde, C. and Taggart, P. (2024). The PopuList: a database of populist, far-left, and far-right parties using expert-informed qualitative comparative classification (EiQCC). British Journal of Political Science, 54(3), 969–978. https://doi.org/10.1017/S0007123423000431

Van Kessel, S. (2015). Populist Parties in Europe: Agents of Discontent? Palgrave Macmillan, London. https://doi.org/10.1057/9781137414113

# 8. Appendix A. The input–output decomposition

Gross output equals intermediate use plus final demand:

$$X = AX + Y \tag{A1}$$

where X is the GN × 1 vector of gross output across all countries and industries, A is the GN × GN matrix of intermediate input coefficients whose element a_sr gives inputs sourced from country s per unit of output in country r, Y is the GN × 1 vector of final demand, G is the number of countries and N the number of industries per country. Rearranging and solving for gross output gives the Leontief system:

$$X = (I - A)^{-1}Y = BY \tag{A2}$$

where I is the identity matrix and B the Leontief inverse. Its element b_sr is the total gross output required from country s, directly and through all indirect upstream rounds, to produce one unit of final demand in country r. The inverse is equivalent to an infinite series:

$$B = I + A + A^2 + A^3 + \ldots \tag{A3}$$

The first term represents direct production, A the first round of input requirements, $A^2$ the inputs needed to produce those inputs, and so on. The Leontief inverse sums the entire chain simultaneously.

Let V be the diagonal matrix of value-added coefficients, each element the share of value added in gross output for country s:

$$v_s = \frac{VA_s}{X_s} \tag{A4}$$

Multiplying V by the Leontief inverse produces the VB matrix, whose element (VB)_sr is the value added generated in country s, required directly and indirectly, to produce one unit of gross output in country r:

$$VB = V \cdot (I - A)^{-1} \tag{A5}$$

Koopman, Wang and Wei (2014) establish the column-sum property of VB. Because all value in any unit of production must originate somewhere, and because value added plus intermediate inputs must equal gross output, the sum along each column equals unity:

$$\sum_s (VB)_{sr} = 1 \qquad \text{for all } r \tag{A6}$$

This property guarantees that every unit of production is fully attributed to domestic or foreign value-added origins with no residual, which is what makes the shares defined in Section 3.1 exhaustive by construction.