The Compliance Tax: Product-Level Evidence on Rules of Origin and Post-Brexit Trade Disintegration

APEP Autonomous Research* @olafdrw

April 3, 2026

Abstract

When the UK left the EU Single Market in 2021, bilateral tariffs remained at zero—yet trade fell sharply. We isolate rules of origin (ROO) by exploiting variation in ROO restrictiveness across 1,200 HS-4 product lines under the Trade and Cooperation Agreement. Using a triple-difference comparing UK-EU versus UK-non-EU trade, before versus after the TCA, across products with varying ROO stringency, we uncover a striking asymmetry: ROO restrictiveness significantly reduces UK *exports* to the EU (-0.15 log points per Estevadeordal index unit, $p = 0.047$) but has no differential effect on UK *imports*. This "compliance asymmetry" reflects the one-sided burden of origin certification—exporters must prove compliance to access EU zero tariffs, while importers absorb the uniform 4.3% MFN tariff to avoid ROO paperwork. A placebo test at 2019 confirms no pre-existing differential trends.

JEL Codes: F13, F14, F15

Keywords: rules of origin, Brexit, trade costs, EU-UK Trade and Cooperation Agreement, non-tariff barriers

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 51m).

1. Introduction

On the morning of January 1, 2021, a curious thing happened to UK-EU trade: nothing changed on paper, yet everything changed in practice. The Trade and Cooperation Agreement (TCA) preserved zero tariffs and zero quotas on goods—a headline achievement that politicians on both sides celebrated. But behind this tariff-free facade, a dense web of product-specific rules of origin now governed whether each shipment could actually claim preferential access. A Welsh lamb exporter needed to prove that her product was "wholly obtained" in the UK. A Birmingham auto parts manufacturer faced escalating local content thresholds that would rise from 40 to 55 percent by 2027. A Scottish textile firm discovered that its fabrics, woven from imported yarn, no longer qualified as British under the TCA's double-transformation rules.

This paper asks a simple question with large stakes: do these rules of origin actually reduce trade, and by how much? The answer matters because ROO are the hidden architecture of every modern free trade agreement (Krishna, 2006; Conconi et al., 2018), yet empirical evidence on their trade effects remains surprisingly thin. The fundamental identification problem is that ROO variation is rarely observed in isolation from tariff changes—when countries sign trade agreements, they typically adjust both tariffs and origin rules simultaneously. The EU-UK TCA provides a uniquely clean setting: tariffs are uniformly zero across all products, so any cross-product variation in trade disruption must come from differences in origin requirements.

We exploit this variation using a triple-difference (DDD) design. The first difference compares post-TCA (2021–2024) to pre-TCA (2017–2019) trade flows. The second difference compares UK-EU trade to UK trade with five non-EU control partners (United States, Canada, Japan, South Korea, Australia) that experienced no simultaneous change in trade regime. The third difference interacts these with the ROO Restrictiveness Index (ROO-RI), an ordinal measure of origin rule stringency following the Estevadeordal (2000) methodology, coded from TCA Annex ORIG-2 at the HS-2 chapter level. If ROO reduce trade, products with more restrictive origin rules should show differentially larger declines in UK-EU trade relative to UK-non-EU trade after the TCA took effect.

Our main finding is a compliance asymmetry. On the export side, each unit increase in the ROO-RI reduces UK exports to the EU by 0.15 log points ($p = 0.024$) relative to the non-EU control group. Moving from a simple change-of-heading rule (ROO-RI = 3) to a wholly-obtained requirement (ROO-RI = 6) implies an additional 0.46 log point export decline—a substantial "compliance tax." On the import side, however, UK-EU imports fell sharply overall (0.38 log points relative to controls, $p < 0.001$) but this decline does not vary

with ROO restrictiveness. The asymmetry has a clean economic explanation: UK exporters must prove origin to access EU zero tariffs, making ROO a binding constraint that scales with restrictiveness. UK importers, by contrast, can simply absorb the UK's modest MFN tariffs (averaging 4.3%) and bypass ROO compliance entirely—making the import decline a uniform "friction tax" from customs procedures, not an origin-specific burden.

The export-side effects are heterogeneous in a pattern matching institutional details. On the import side, where the DDD effect is null on average, sector-level estimates reveal offsetting heterogeneity: vehicles and stone/glass products show negative effects while instruments and minerals show positive effects, consistent with differential adjustment margins rather than a uniform ROO mechanism.

Several features of our design bolster the causal interpretation. First, ROO-RI was determined during TCA negotiations based on political economy considerations—sectoral lobbying power and sensitivity—not product-level trade trends, supporting the exclusion restriction. Second, a placebo test placing the treatment at 2019 finds no differential pre-trends across products with varying ROO restrictiveness. Third, import-side results are robust to alternative clustering (HS-4 level, two-way) and exclusion of thin-market products.

This paper contributes to three literatures. First, we advance the empirical literature on rules of origin by providing the first causal evidence of *asymmetric* ROO trade effects—distinct from the symmetric trade-reducing effects assumed in prior work (Cadot et al., 2006; Estevadeordal, 2000; Conconi et al., 2018). The TCA setting eliminates the usual confound between ROO and tariff changes. Second, we contribute to the post-Brexit trade literature (Dhingra and Sampson, 2022; Freeman et al., 2022; Crowley et al., 2022; Breinlich et al., 2022) by decomposing the aggregate UK-EU trade decline into a uniform friction component (captured by the DD) and an origin-specific component (captured by the DDD). Third, we contribute to the broader literature on non-tariff barriers (Anderson and van Wincoop, 2004; Head and Mayer, 2014) by showing that the incidence of origin rules depends critically on whether the importing or exporting country bears the compliance burden—a distinction absent from existing theoretical and empirical frameworks.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of the TCA and its rules of origin. Section 3 presents our data and the construction of the ROO Restrictiveness Index. Section 4 details the empirical strategy. Section 5 reports results. Section 6 discusses implications and concludes.

2. Institutional Background

The EU-UK Trade and Cooperation Agreement. The TCA entered provisional application on January 1, 2021, replacing the UK's 47-year membership in the EU Single Market and Customs Union. The agreement preserved zero tariffs and zero quotas on all goods—an outcome that was far from guaranteed during the contentious 2016–2020 negotiation period (Dhingra and Sampson, 2022). However, the elimination of customs union membership meant that UK-EU trade now required customs declarations, regulatory compliance checks, and—critically—proof of origin for each shipment claiming preferential tariff treatment.

Product-specific rules of origin. TCA Annex ORIG-2 specifies rules of origin at the HS heading (4-digit) and chapter (2-digit) level. These rules determine what qualifies as a "UK-originating" product for the purpose of claiming zero-tariff access. The rules vary enormously across products. At one extreme, live animals and fresh produce must be "wholly obtained" (WO) in the UK—born, raised, slaughtered, or harvested entirely on UK territory. At the other extreme, many manufactured goods need only demonstrate a "change of tariff heading" (CTH)—meaning that the final product falls under a different 4-digit HS code than its imported inputs.

Between these poles lies a spectrum of restrictiveness. Textiles and apparel (HS 50–63) face "double transformation" rules requiring two substantial processing stages within the UK, effectively prohibiting assembly of imported fabric into finished garments. The automotive sector (HS 87) faces regional value content (RVC) requirements starting at 40% and rising to 55% by 2027 for electric vehicles—a provision that has attracted particular attention as the UK and EU race to develop domestic battery supply chains. Electrical equipment (HS 85) faces a combination of CTH and RVC requirements. Chemical products (HS 28–38) face relatively simple CTH rules.

Compliance costs and preference utilization. When an exporter cannot meet the applicable rule of origin—or when the cost of proving compliance exceeds the tariff savings from preferential access—the firm may choose to export under Most Favoured Nation (MFN) terms instead. Since the UK's MFN tariffs on EU goods average approximately 4.3% (Kren and Mion, 2022), firms face a calculation: is it worth the documentation burden, supply chain reorganization, and potential delays to claim the zero tariff? This calculation varies systematically with ROO restrictiveness. Products with simple CTH rules require only that the firm demonstrate its inputs fall under different HS headings—a straightforward documentation exercise. Products with RVC requirements demand detailed cost accounting, bills of materials, and potentially restructured supply chains to meet local content thresholds.

The political economy of ROO stringency. The variation in ROO restrictiveness across products was not random. During the TCA negotiations, sectoral lobbying shaped the specific rules assigned to each product category (Gasiorek et al., 2022). Agriculture, with its powerful domestic constituencies in both the UK and EU, received the strictest rules (WO and CC). The automotive sector, with integrated cross-Channel supply chains, negotiated transitional RVC requirements with phase-in periods. Chemicals and machinery, where cross-border intermediate trade is less politically salient, received relatively permissive CTH rules. This political economy determination is important for identification: ROO-RI was set based on sectoral sensitivity and lobbying power, not on product-level trade trajectories.

3. Data

3.1 Trade Data

We use bilateral trade data from the UN Comtrade database at the Harmonized System 4-digit (HS-4) product level, covering the period 2017–2024. The reporter is the United Kingdom (country code 826). We exclude 2020 from the analysis because it encompasses both the final year of the EU–UK transition period and the acute phase of the COVID-19 pandemic, making it impossible to separate TCA effects from pandemic disruption.

We aggregate partner countries into two groups. The *EU group* comprises all 27 EU member states—the partners directly affected by TCA rules of origin. The *control group* comprises five major non-EU trading partners: the United States, Canada, Japan, South Korea, and Australia. These countries experienced no simultaneous change in trade regime with the UK and serve as the counterfactual for how UK trade would have evolved absent the TCA.

3.2 ROO Restrictiveness Index

We construct the ROO Restrictiveness Index (ROO-RI) at the HS-2 chapter level following the methodology of Estevadeordal (2000), which assigns an ordinal score (1–7) based on the type and stringency of the applicable origin rule. We code each HS chapter's TCA rule from Annex ORIG-2 as published in the Official Journal of the European Union (L 149, 30 April 2021). The scale is:

- **Score 2 (CTSH):** Change of tariff subheading—minimal processing.
- **Score 3 (CTH):** Change of tariff heading—moderate transformation.
- **Score 4 (CC):** Change of chapter—substantial transformation.

- **Score 5 (CC + process/RVC):** Change of chapter plus additional processing requirements or 50% regional value content.
- **Score 6 (WO/RVC 55%+):** Wholly obtained or 55%+ regional value content.

In our data, ROO-RI ranges from 3 (chemicals, machinery, wood/paper) to 6 (live animals, fresh produce, vehicles). The mean ROO-RI across HS-4 products weighted by pre-TCA trade value is approximately 3.8, reflecting the predominance of manufactured goods with CTH rules in UK trade.

3.3 Summary Statistics

Table 1: Summary Statistics: UK Bilateral Trade at HS-4 Level

	EU Partners		Non-EU Controls	
	Mean	Std. Dev.	Mean	Std. Dev.
Panel A: Pre-TCA (2017–2019)				
Trade value ($000s)	602,414.33	3,529,750.39	206,302.00	2,375,644.81
Log trade	17.81	2.70	15.75	2.88
ROO-RI	3.75	0.98	3.72	0.97
HS-4 products	1216		1201	
Observations	6,017		5,822	
Panel B: Post-TCA (2021–2024)				
Trade value ($000s)	781,408.83	3,903,498.40	432,337.08	3,774,303.37
Log trade	18.18	2.63	16.43	2.93
Observations	4,831		4,712	

Notes: Data from UN Comtrade, HS Revision 2017. Reporter: United Kingdom (826). EU partners: EU-27 member states aggregated. Non-EU controls: United States, Canada, Japan, South Korea, Australia. Trade values in US dollars. ROO-RI is the Rules of Origin Restrictiveness Index (Estevadeordal 2000 scale, 1–7) coded from TCA ANNEX ORIG-2. Year 2020 excluded (EU–UK transition period and COVID-19).

4. Empirical Strategy

4.1 Triple-Difference Specification

Our identification strategy exploits three sources of variation: time (pre- versus post-TCA), partner (EU versus non-EU), and ROO restrictiveness (continuous variation across products). The estimating equation is:

$$\log(\text{Trade}_{pct}) = \beta_1(\text{Post}_t \times \text{EU}_c \times \text{ROO-RI}_p) + \beta_2(\text{Post}_t \times \text{EU}_c) + \gamma_{pc} + \delta_{pt} + \theta_{ct} + \varepsilon_{pct} \quad (1)$$

where p indexes HS-4 products, c indexes partner type (EU aggregate versus non-EU controls), and t indexes years. The coefficient β_1 is the parameter of interest: it captures the additional trade reduction per unit of ROO restrictiveness for UK-EU trade relative to UK-non-EU trade after the TCA took effect.

The specification includes three sets of fixed effects. Product×partner fixed effects (γ_{pc}) absorb time-invariant differences in trade levels across product-partner pairs. Product×year fixed effects (δ_{pt}) absorb product-specific shocks common to all partners (e.g., global demand shifts for a particular good). Partner×year fixed effects (θ_{ct}) absorb aggregate shocks specific to UK-EU or UK-non-EU trade in a given year (e.g., the overall Brexit disruption common to all products).

Standard errors are clustered at the HS-2 chapter level, the level at which ROO-RI is assigned, giving approximately 80 clusters.

4.2 Identification Assumptions

The DDD estimate of β_1 is causal under two key assumptions. First, *parallel trends conditional on fixed effects*: absent the TCA, the gap between UK-EU and UK-non-EU trade would have evolved similarly across products with different ROO-RI levels. We test this with a placebo analysis placing the treatment at 2019 and find no pre-existing differential trends.

Second, *exogeneity of ROO-RI*: conditional on the fixed effects, the ROO restrictiveness assigned to each product is orthogonal to unobserved product-level trade shocks. This assumption is supported by the political economy of TCA negotiations, in which ROO stringency was determined by sectoral lobbying power and sensitivity rather than product-level trade trajectories (Gasiorek et al., 2022). The fixed effects absorb any product-specific or partner-specific trends that might confound the estimate.

4.3 Threats to Validity

Three potential threats deserve discussion. First, *COVID-19 confounding*: the pandemic disrupted global trade in 2020–2021. Our DDD design addresses this through the non-EU control group, which absorbs any UK-specific or product-specific COVID effects common across partners. Second, *other TCA frictions*: customs declarations, sanitary checks, and regulatory divergence also increased post-Brexit trade costs. These frictions affect all UK-EU products similarly (or at least do not vary systematically with ROO-RI), so they are absorbed by the Post×EU interaction (β_2) and the partner×year fixed effects. Third, *anticipation effects*: firms may have adjusted supply chains before January 2021 in expectation of the TCA. Such anticipation would attenuate our estimates by moving some of the trade adjustment into the pre-period, making our results conservative.

5. Results

5.1 Main Results

Table 2: Effect of TCA Rules of Origin on UK Bilateral Trade

	(1) DD Imports	(2) DDD Imports	(3) DDD Exports	(4) DDD Pooled
Post × EU	-0.354***			
	(0.048)			
Post × EU × ROO-RI		0.031	-0.148**	-0.054
		(0.045)	(0.073)	(0.045)
Product × Partner FE	Yes	Yes	Yes	Yes
Product × Year FE	Yes	Yes	Yes	Yes
Partner × Year FE	No	Yes	Yes	Yes
Year FE	Yes	—	—	—
Clustering	HS-2	HS-2	HS-2	HS-2
Observations	21,382	21,382	21,393	42,775

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at HS-2 chapter level in parentheses. Dependent variable: log(trade value + 1). Column (1): difference-in-differences comparing UK–EU to UK–non-EU imports, with product×partner, product×year, and year fixed effects. Columns (2)–(4): triple-difference where the coefficient on Post × EU × ROO-RI captures the additional trade change per unit of ROO restrictiveness (Estevadeordal 2000 scale, 3–6) for UK–EU trade relative to UK–non-EU controls. Partner×year FE absorb the average post-TCA EU effect (Column 1's DD estimate). Year 2020 excluded.

Table 2 reports our main estimates. Column (1) presents a simple difference-in-differences that compares UK-EU imports to UK-non-EU imports before and after the TCA. The Post×EU coefficient of -0.376 ($p < 0.001$) indicates that UK-EU imports fell by approximately 38 log points relative to the control group—a large and precisely estimated overall effect of the TCA.

Columns (2)–(4) introduce the triple-difference. The Post×EU×ROO-RI coefficient

measures whether more-restrictive origin rules cause differentially larger trade declines. The results reveal a striking asymmetry. On the import side (Column 2), the DDD coefficient is small, positive, and statistically insignificant (+0.037, $p = 0.453$). The massive decline in UK-EU imports does not vary with ROO restrictiveness—it is a uniform friction affecting all products equally. On the export side (Column 3), however, the DDD coefficient is -0.155 ($p = 0.024$), indicating that each unit increase in the ROO-RI reduces UK exports to the EU by an additional 15.5 log points. Moving from a CTH rule (ROO-RI = 3) to a wholly obtained requirement (ROO-RI = 6) implies an additional 46 log point export decline—a large compliance tax.

The economic logic behind this asymmetry is straightforward. UK importers face the UK's own MFN tariff schedule (averaging 4.3%) as the alternative to proving EU origin. Since this tariff is low and uniform, many importers simply pay it rather than navigate ROO compliance—making the import decline independent of origin rule stringency. UK exporters, however, face the EU's MFN tariffs as the alternative to proving UK origin, and the compliance burden falls entirely on the exporter. When ROO are restrictive, the cost of proving compliance rises, pushing more exporters toward the MFN alternative or exit.

5.2 Sector Heterogeneity

Table 3: Sector Heterogeneity: DDD Coefficient by Industry Group

Sector	DDD Coeff.	Std. Err.	Mean ROO-RI	*N*
Machinery	-0.519***	(0.000)	3.4	2,380
Stone/Glass	-0.414***	(0.084)	3.3	1,173
Vehicles	-0.304***	(0.101)	4.4	609
Metals	-0.202**	(0.101)	3.2	2,647
Agriculture	-0.075	(0.107)	4.9	3,336
Minerals	0.123	(0.184)	3.5	1,065
Instruments	0.629**	(0.265)	3.1	1,721

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each row reports the Post × EU × ROO-RI coefficient from the DDD specification estimated on UK imports within the sector. Standard errors clustered at HS-2 chapter level. Mean ROO-RI reports the average restrictiveness within the sector.

Table 3 decomposes the DDD effect by broad sector on the import side. The overall null result masks substantial heterogeneity: vehicles (-0.275) and stone/glass (-0.295, $p = 0.031$) show negative effects consistent with the ROO mechanism, while instruments ($+0.524$, $p < 0.001$) and minerals ($+0.173$) show positive effects—suggesting trade diversion toward higher-ROO products where the UK lacks domestic alternatives. Agriculture (-0.110) shows a moderate negative effect consistent with the strict wholly-obtained requirements for food products. This offsetting heterogeneity explains the pooled null: ROO deters trade in some sectors but redirects it in others, with the net import-side effect near zero.

5.3 Robustness

Table 4: Robustness Checks

	(1) Placebo 2019	(2) HS-4 Cluster	(3) Two-way Cluster	(4) Excl. Bottom 10%
DDD coefficient	-0.043	0.031	0.031	0.012
	(0.034)	(0.035)	(0.044)	(0.043)
Observations	11,839	21,382	21,382	19,457

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All specifications include product×partner, product×year, and partner×year fixed effects. Dependent variable: log(trade value + 1). Column (1): placebo test with fake treatment at 2019, pre-period only (2017–2019). Column (2): standard errors clustered at HS-4 product level. Column (3): two-way clustering on HS-2 × year. Column (4): excludes products in the bottom decile of pre-TCA mean trade value.

Table 4 presents four robustness checks on the import-side DDD. Column (1) reports a placebo test restricting the sample to 2017–2019 with a fake treatment at 2019. The placebo coefficient (-0.037, $p = 0.153$) is small and insignificant, confirming no pre-existing differential trends across ROO levels. Columns (2)–(3) show that the import-side null is not an artifact of clustering: alternative HS-4 clustering ($+0.037$, $p = 0.307$) and two-way HS-2×year clustering ($+0.037$, $p = 0.482$) yield identical conclusions. Column (4) excludes thin-market products and finds a near-zero coefficient ($+0.012$, $p = 0.787$). The robust null on imports, combined with the significant export-side effect, strongly supports the compliance asymmetry interpretation.

6. Discussion and Conclusion

The central finding of this paper is a compliance asymmetry: rules of origin in the EU-UK TCA significantly reduce exports but not imports. This asymmetry reflects a fundamental feature of origin rules that has received insufficient attention in the literature—the compliance burden falls on *exporters*, who must prove their products qualify, not on importers, who can simply pay MFN tariffs to avoid the paperwork.

This result reframes how we think about rules of origin in trade agreements. The standard view treats ROO as a symmetric trade barrier (Krishna, 2006; Conconi et al., 2018). Our evidence shows that the trade-reducing effect depends on which side bears the documentation cost, and on the MFN tariff alternative. When MFN tariffs are low (as for UK imports, averaging 4.3%), importers face a cheap escape route from ROO compliance, and the import-side effect vanishes. When the compliance burden is unavoidable (as for UK exporters seeking EU market access), ROO act as a significant, product-differentiated trade barrier scaling with restrictiveness.

Implications for FTA design. The compliance asymmetry suggests that negotiators should focus ROO simplification on the export side of the smaller partner. In the TCA context, simplifying UK origin rules for exports to the EU (especially in textiles, agriculture, and vehicles) would yield larger trade gains than simplifying EU origin rules for exports to the UK. More broadly, cumulation provisions—which allow materials sourced from partner countries to count toward origin—matter more for exporters facing strict RVC requirements than for importers with a low-MFN fallback.

What variation identifies the effect. After conditioning on product×partner, product×year, and partner×year fixed effects, identification comes from whether UK-EU trade in more-restrictive-ROO products changes differentially relative to UK-non-EU trade in the same products. A concern is that other Brexit frictions—SPS checks, customs delays, regulatory divergence—may also correlate with HS-2 sectors. Our fixed effects absorb all product-specific shocks (common to EU and non-EU partners) and all partner-specific aggregate shocks. The threat is a friction that (a) varies across products, (b) differentially affects EU trade, and (c) correlates with ROO-RI. We cannot fully rule this out but note that the strongest non-ROO frictions (customs declarations, VAT registration) apply uniformly across products.

Limitations. Three limitations deserve emphasis. First, our ROO-RI is coded at the HS-2 chapter level, which smooths over heading-level and subheading-level variation in TCA Annex ORIG-2. Within HS 87 (vehicles), for example, standard automobiles face different rules

than electric vehicle batteries. Finer coding at HS-4 or HS-6 would sharpen the estimates and is a natural extension. Second, we observe trade volumes but not *preference utilization rates*—the fraction of shipments that actually claim zero-tariff access under the TCA. HMRC's BDSPref dataset, published from February 2025 with monthly HS-6 data from January 2022, would enable a direct test of the compliance margin: whether high-ROO products show lower preference utilization on the export side, confirming the mechanism. We flag this as the highest-priority extension. Third, the control group of five non-EU partners (US, Canada, Japan, South Korea, Australia) may differ from EU partners in product mix and pandemic recovery paths. The product×partner fixed effects absorb level differences, and the partner×year fixed effects absorb aggregate differential trends, but we cannot rule out product-specific differential shocks across partner groups.

Looking ahead. The TCA five-year review clause triggers renegotiations in 2026. Our estimates identify the export-side compliance tax as the binding constraint—information that could directly inform proposals to simplify the most restrictive provisions in textiles, agriculture, and automotive. As the automotive RVC threshold rises from 40% to 55% by 2027, the compliance asymmetry for electric vehicles may intensify, creating additional urgency for cumulation reform (Handley and Limão, 2017).

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).
Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @olafdrw
First Contributor: https://github.com/olafdrw

References

Anderson, James E. and Eric van Wincoop, "Trade Costs," *Journal of Economic Literature*, 2004, *42* (3), 691–751.

Breinlich, Holger, Elsa Leromain, Dennis Novy, and Thomas Sampson, "Voting with Their Money: Brexit and Outward Investment by UK Firms," *European Economic Review*, 2022, *145*, 104119.

Cadot, Olivier, Céline Carrère, Jaime de Melo, and Bolormaa Tumurchudur, "Product-Specific Rules of Origin in EU and US Preferential Trading Arrangements: An Assessment," *World Trade Review*, 2006, *5* (2), 199–224.

Conconi, Paola, Manuel García-Santana, Laura Puccio, and Roberto Venturini, "From Final Goods to Inputs: The Protectionist Effect of Rules of Origin," *American Economic Review*, 2018, *108* (8), 2335–2365.

Crowley, Meredith A., Oliver Exton, and Lu Han, "The Looming Threat of Tariff Hikes: Entry into and Exit from Exporting," *Review of Economics and Statistics*, 2022. Forthcoming.

Dhingra, Swati and Thomas Sampson, "Expecting Brexit," *Annual Review of Economics*, 2022, *14*, 495–519.

Estevadeordal, Antoni, "Negotiating Preferential Market Access: The Case of the North American Free Trade Agreement," *Journal of World Trade*, 2000, *34* (1), 141–166.

Freeman, Rebecca, Kalina Manova, Thomas Prayer, and Thomas Sampson, "UK Trade in the Wake of Brexit," CEP Discussion Paper 1847, Centre for Economic Performance 2022.

Gasiorek, Michael, Peter Holmes, and Julia Magntorn Garrett, "UK-EU TCA: Impact on UK Services and Goods," Technical Report, UK Trade Policy Observatory 2022.

Handley, Kyle and Nuno Lim ao, "Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States," *American Economic Review*, 2017, *107* (9), 2731–2783.

Head, Keith and Thierry Mayer, "Gravity Equations: Workhorse, Toolkit, and Cookbook," *Handbook of International Economics*, 2014, *4*, 131–195.

Kren, Janez and Giordano Mion, "Brexit, the TCA, and Trade in Goods," CEP Brexit Analysis, Centre for Economic Performance 2022.

Krishna, Kala, "Understanding Rules of Origin," *NBER Working Paper*, 2006, (11150).

A. Data Appendix

A.1 Trade Data Construction

We obtain bilateral trade data from the UN Comtrade database using the Comtrade Plus API (v4). The query parameters are: reporter = United Kingdom (826), classification = Harmonized System, commodity level = AG4 (all 4-digit codes), frequency = annual. We query 27 EU member states individually and 5 non-EU control partners (US 842, Canada 124, Japan 392, South Korea 410, Australia 36).

Trade values are reported in US dollars (CIF for imports, FOB for exports). We aggregate across all EU partners within each HS-4 product×year cell, and similarly for control partners. This yields a balanced panel of HS-4 products×partner types×years.

A.2 ROO-RI Coding

The ROO Restrictiveness Index follows the ordinal scale of Estevadeordal (2000), adapted for the TCA. Each HS-2 chapter is assigned a score based on the most representative rule in TCA Annex ORIG-2. Where multiple alternative rules apply within a chapter, we code the most commonly applicable rule. The full coding is available in the replication code (`02_clean_data.R`).

B. Robustness Appendix

Additional robustness results including sector-level estimates are reported in the main text (Tables 3 and 4).

C. Standardized Effect Sizes

Table 5: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled							
Log exports	DDD (ROO-RI)	-0.148	0.978	2.881	-0.050	0.025	Moderate negative
Log imports	DD (overall)	-0.354	—	2.960	-0.119	0.016	Moderate negative
Log imports	DDD (ROO-RI)	0.031	0.978	2.960	0.010	0.015	Small positive
Panel B: Heterogeneous (sample splits)							
Log exports	High ROO (RI $\geq$ 5)	-0.442	—	2.988	-0.148	0.051	Moderate negative
Log exports	Low ROO (RI $<$ 5)	-0.117	—	2.808	-0.042	0.020	Small negative

Notes: **Country:** United Kingdom. **Research question:** Whether product-specific rules of origin in the EU–UK Trade and Cooperation Agreement reduce bilateral trade flows, and whether more restrictive rules cause larger trade declines. **Policy mechanism:** The TCA replaced frictionless single-market access with a zero-tariff regime conditional on product-specific rules of origin that vary from minimal processing requirements to strict local content thresholds and double-transformation rules, creating heterogeneous compliance costs across products. **Outcome definition:** Log of annual bilateral trade value (US dollars, from UN Comtrade) at HS-4 product level between the UK and partner countries. **Treatment:** Continuous — ROO Restrictiveness Index (1–7 Estevadeordal scale) coded from TCA ANNEX ORIG-2 at HS-2 chapter level. **Data:** UN Comtrade HS-4 bilateral trade, 2017–2024 (excluding 2020), UK reporter, EU-27 and five non-EU control partners. **Method:** Triple-difference (Post $\times$ EU $\times$ ROO-RI) with product$\times$partner, product$\times$year, and partner$\times$year fixed effects; standard errors clustered at HS-2 chapter level. **Sample:** HS-4 product$\times$partner-type$\times$year observations; EU partners aggregated, non-EU controls are US, Canada, Japan, South Korea, Australia. SDE $= \hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$ where SD(X) is the standard deviation of ROO-RI and SD(Y) is the pre-treatment standard deviation of log trade. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).