CITP
CENTRE FOR INCLUSIVE TRADE POLICY

The trade elasticity from tariff-based regressions: what do we measure?

Nicolò Tamberi

February 2026

Centre for Inclusive Trade Policy
Working Paper No.031

Economic and Social Research Council

Centre for Inclusive Trade Policy
https://citp.ac.uk/
info@citp.ac.uk

Established in 2022, the Centre for Inclusive Trade Policy (CITP) is the first research centre dedicated to trade policy to be funded by the Economic and Social Research Council. As a centre of excellence for innovative research on trade policy and its inclusiveness, we aim to equip the UK with the capability to formulate and implement a trade policy tailored to the needs of the whole of the UK, while recognising the importance of the multilateral trading system and the UK's role within it. The CITP is funded by the Economic and Social Research Council [grant number ES/W002434/1]

This Working Paper is issued under the auspices of the Centre's research programme. Any opinions expressed here are those of the author(s) and not those of the Centre for Inclusive Trade Policy. Research disseminated by CITP may include views on policy, but the Centre itself takes no institutional policy positions.

These Working Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character. The author(s) asserts their moral right to be identified as the author of this publication. For online use, we ask readers to link to the webpage for the resource on the Centre's website.

Abstract

Most empirical estimates of the trade elasticity exploit tariff variation generated by Preferential Trade Agreements (PTAs). Because firms only partly utilise preferential tariffs, standard regressions identify an eligible-tariff elasticity rather than the structural elasticity with respect to tariffs actually paid that is required for quantitative trade models. I model PTA tariff cuts as a treatment with partial compliance and estimate the effective-tariff elasticity using an instrumental variable approach that explicitly accounts for preference utilisation. The effective-tariff elasticity lies between 10 and 12, roughly three times larger than eligible-tariff estimates around 3.5. I show that the eligible-tariff elasticity varies with the tariff preference utilisation of the estimation sample and is therefore not portable across policy environments, while the effective-tariff elasticity identifies a policy-invariant elasticity that can be used for quantitative modelling.

©: Nicolò Tamberi

Suggested citation

N, Tamberi; (2026) The trade elasticity from tariff-based regressions: what do we measure?, Centre for Inclusive Trade Policy, Working Paper 031

Non-Technical Summary

A large share of empirical research on international trade relies on variation in import tariffs to estimate how strongly trade responds to changes in trade costs. Crucially, most of this tariff variation is generated by Preferential Trade Agreements, which allow importers to charge lower tariffs to certain partners. Because these agreements specify different tariff rates for different exporters, they create rich variation that economists routinely exploit to estimate the "trade elasticity" – a key parameter used in nearly all quantitative trade models.

However, a central but often overlooked fact is that firms do not automatically receive these preferential tariffs. To claim them, they must comply with rules of origin and other administrative requirements that can be costly or difficult. As a result, many eligible exporters continue paying the higher Most Favoured Nation (MFN) tariff instead of using the preferential rate. The share of imports that actually enters under the preferential tariff – the preference utilisation rate – is often far below 100%.

This paper demonstrates that ignoring these low utilisation rates leads to a fundamental mismeasurement of the trade elasticity. Standard regressions of imports on statutory preferential tariffs identify the responsiveness of trade to eligible tariffs rather than to the effective tariffs actually paid. In statistical terms, this is a classic partial compliance problem: tariff preferences are the treatment, but only some eligible firms comply by using them. The estimated coefficient therefore corresponds to an intent to treat effect, not the true treatment effect.

This paper shows that the commonly reported "tariff elasticity" from the literature is an eligible tariff elasticity, equal to the true structural elasticity multiplied by the average preference utilisation rate in the estimation sample. Because preference utilisation rates typically lie well below one, traditional estimates are attenuated. Moreover, these eligible tariff elasticities are not policy invariant: they depend on the specific utilisation patterns in the data – which vary widely across partners and agreements – and therefore cannot be validly applied to other policy environments, such as MFN tariff reforms or customs unions, where compliance is effectively complete.

To recover the correct effective tariff elasticity, the paper applies an instrumental variables approach inspired by the Local Average Treatment Effect (LATE) literature. Tariff reductions implied by EU Preferential Trade Agreements serve as the treatment assignment, while the tariff reductions actually realised – statutory cuts multiplied by the observed preference utilisation rate – represent the treatment received. Using this strategy on four major EU agreements (with Korea, Canada, Japan, and Singapore), the paper finds that the effective-tariff elasticity lies between 10 and 12 – around three times larger than eligible-tariff elasticity estimates of 3-4.

These findings have major implications for quantitative trade modelling. Because most models currently use eligible tariff elasticities, they effectively treat partial utilisation as if it reduced the responsiveness of trade rather than the size of the tariff change. This leads to systematic overprediction of the trade effects of FTAs. A general equilibrium illustration using the EU-Japan agreement shows that standard modelling approaches can overstate predicted trade increases by up to a factor of two.

The paper concludes that researchers and policymakers should treat partial utilisation as a feature of the tariff change – not as a modification of the elasticity – and that future empirical work should incorporate preference utilisation rates directly. Building comprehensive datasets on utilisation rates will be essential to ensure that trade policy is evaluated using structurally meaningful parameters.

The trade elasticity from tariff-based regressions: what do we measure?

Nicolò Tamberi*

19th February 2026

Abstract

Most empirical estimates of the trade elasticity exploit tariff variation generated by Preferential Trade Agreements (PTAs). Because firms only partly utilize preferential tariffs, standard regressions identify an eligible-tariff elasticity rather than the structural elasticity with respect to tariffs actually paid that is required for quantitative trade models. I model PTA tariff cuts as a treatment with partial compliance and estimate the effective-tariff elasticity using an instrumental variable approach that explicitly accounts for preference utilization. The effective-tariff elasticity lies between 10 and 12, roughly three times larger than eligible-tariff estimates around 3.5. I show that the eligible-tariff elasticity varies with the tariff preference utilization of the estimation sample and is therefore not portable across policy environments, while the effective-tariff elasticity identifies a policy-invariant elasticity that can be used for quantitative modelling.

1 Introduction

One of the most fundamental parameters in international economics is the elasticity of imports to variable trade costs, often referred to as the trade elasticity. This parameter governs the response of trade flows to changes in trade policy and it is central to the evaluation of GDP and welfare effects in quantitative trade models.

Empirically, a common approach to estimating the trade elasticity is to exploit variation in import tariffs across exporting countries and products. In practice, much of this identifying variation is generated by tariff preferences granted under Preferential Trade Agreements (PTAs). However, this widely used approach overlooks a key feature of PTAs: tariff preferences are often only partially utilized.

*University of Sussex. This research was supported by the ESRC Centre for Inclusive Trade Policy, grant number ES/W002434/1.

This paper revisits the estimation of the trade elasticity from import tariffs by explicitly accounting for the partial uptake of PTA tariff preferences. Interpreting tariff preferences as treatment, their partial uptake put us in a situation of partial compliance with treatment assignment: some of the firms eligible for preferential tariffs do not use them and trade under non-preferential terms. I show that standard regressions of imports on tariffs identify an eligible-tariff elasticity – the response of imports to statutory tariff changes – rather than the elasticity with respect to tariffs actually paid, which I refer to as the effective-tariff elasticity. Accounting for preference utilization recovers the effective-tariff elasticity, which is substantially larger than the eligible-tariff elasticity.

Crucially, the eligible-tariff elasticity is not a structural parameter: it is a reduced-form object that combines the structural elasticity with the average preference utilization rate of the estimation sample. This object is therefore non-portable across policy environments and subject to the Lucas critique. As such, it is not the object required as an input in quantitative trade models. By contrast, the effective-tariff elasticity identifies a policy-invariant primitive – the structural elasticity – that can be used in counterfactual analysis.

Conceptually, the setting can be viewed as a problem of treatment with one-sided partial compliance: tariff preferences (the treatment assignment) are only partially utilized by eligible firms. This parallels the logic of the Local Average Treatment Effect/Intent-To-Treat (LATE/ITT) framework in applied micro-econometrics. To address the partial compliance issue, I use an IV strategy that recovers the elasticity of imports with respect to the effective tariff – the analogue of a compliance-corrected LATE. Applying this framework to European Union trade agreements, I find that the effective-tariff elasticity lies in the range of 10-12, compared to values around 3-4 for the eligible-tariff elasticity.

Under the rules of the World Trade Organization, countries apply a common external tariff across all exporters, known as the Most Favoured Nation (MFN) tariff. Importers may then grant lower tariffs to specific exporters through Preferential Trade Agreements, either bilaterally, or unilaterally, as in preference schemes extended by developed countries to developing ones. Access to these preferential tariffs is not automatic: exporters (or importers) must demonstrate that goods satisfy certain criteria, with the most important probably being rules of origin. Complying with these requirements can be costly or infeasible, leading many eligible exporters to not use the low preferential tariffs.

The extent to which tariff preferences are used is measured by the Preference Utilization Rate (PUR), defined as the share of eligible imports that actually use preferential tariffs. Figure 1 reports the average PUR across products of European Union (EU) imports from South Korea, Canada, Japan and Singapore. Even Korea, the country with the highest average PUR, just reaches 40% at the product level. Although a separate literature has documented the determinants of preference utilization, its implications for the estimation of trade elasticities and for quantitative trade models have largely been overlooked.

Figure 1: Preference Utilization Rates of the EU over time

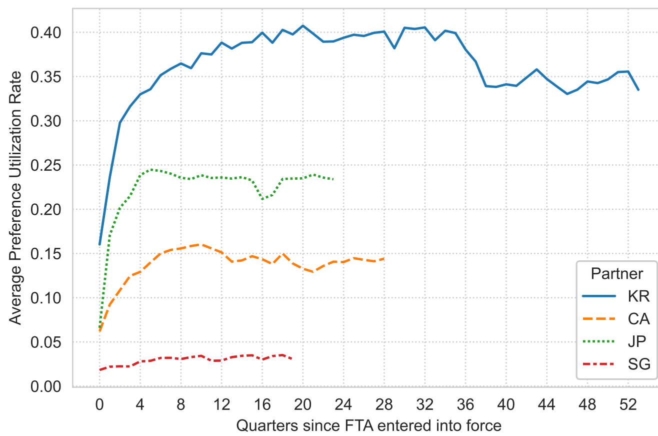

Source: author's elaboration of Eurostat data. The figure show the average preference utilization rate of EU imports from Korea, Canada, Japan and Singapore. The PUR is calculated on products with a preferential tariff below the MFN rate, averaged across CN 8-digit products.

Because tariff preferences are only partially utilized, we are in a one-sided non-compliance setting. Standard regressions of imports on tariffs capture the elasticity of tariff eligibility – effectively the structural trade elasticity multiplied by the average PUR. Interpreting this object as the trade elasticity in quantitative trade models leads to a systematic attenuation, as estimates depend on the in-sample average PUR, which varies between zero and one. This problem arises both with firm- and product-level data.

To address this problem, I model PTA tariff preferences as a continuous treatment subject to partial compliance and estimate the trade elasticity using an instrumental variable approach. Treatment assignment, given by the tariff reduction from MFN to PTA levels, is used as an instrument for received treatment, measured as the effective tariff reduction – the statutory tariff reduction interacted with the PUR. The analysis is implemented using local projections, allowing me to estimate cumulative effects over time. Furthermore, I address the potential endogeneity of tariff preferences granted to one country with the average tariff preference granted to other countries.

The use of treatment assignment as an instrument for treatment received follows the standard LATE/ITT logic in micro-econometrics. The effective preferential tariff faced by firms equals the statutory tariff reduction multiplied by the preference utilization rate (PUR). Because PUR is an endogenous equilibrium object, jointly determined with import flows and related to self-selection of firms into using tariff preferences, it is mechanically correlated with the error term. This makes the effective tariff change endogenous.

If statutory tariff changes were exogenous, one could directly use them as an instrument for the effective tariff change. However, PTA tariff schedules may be subject to political-economy endogeneity: sectors that are more competitive or more sensitive may receive smaller or slower tariff cuts. This creates potential correlation between a country's own statutory tariff changes and unobserved determinants of its import patterns.

To address this, I instrument the effective tariff change using the statutory tariff changes scheduled under other PTAs. These 'external' statutory changes provide plausibly exogenous variation because they are negotiated independently of the focal country's sector-specific shocks. At the same time, they remain highly correlated with the effective tariff change through the common structure of tariff schedules across PTAs. This satisfies the relevance condition while avoiding the political-economy endogeneity that affects the country's own tariff commitments, ensuring that the validity of the instrument.

The empirical results align closely with the theoretical predictions. Ignoring preference utilization yields tariff elasticities between 3 and 4, consistent with much of the existing literature. Once partial compliance is accounted for, the implied elasticity rises to 10-12. Consistent with the LATE interpretation, the eligible-tariff estimate is approximately equal to the effective-tariff elasticity multiplied by the average preference utilization rate.

The goal of this paper is not to rescale coefficients but to clarify the object of interest for structural modelling and policy evaluation. For structural modelling, the parameter of interest is not the eligible-tariff elasticity, which equals the structural elasticity scaled by the preference utilization rate, but the effective-tariff elasticity, which identifies the policy-invariant structural elasticity. I advance two core arguments.

First, the eligible-tariff elasticity is not portable across policy environments. Since it embeds the in-sample preference utilization rate, it will be inappropriate for settings where PUR differs – most notably policies with full compliance, such as customs unions or MFN tariff changes. In contrast, the effective-tariff elasticity corresponds directly to the structural elasticity and is consistent across such environments.

Second, the core implication for structural modelling is that partial preference uptake should enter through the tariff, not through the elasticity. Standard practice that uses the eligible-tariff elasticity implicitly models partial compliance by scaling down the trade elasticity rather than the tariff change itself. However, theory-consistent modelling requires scaling the tariff change, not the elasticity. Because the relationship between tariff changes and trade flows is non-linear in standard CES trade models, rescaling the elasticity rather than the tariff change leads to systematic overpredictions of the trade effects of tariff liberalizations in PTAs. I illustrate this point numerically using the EU-Japan trade agreement, showing how incorporating effective tariff reductions and preference utilization rates materially changes predicted trade effects.

Finally, the issues raised in this paper extend beyond tariff preferences to other policy settings involving partial uptake or conditional treatment. A prominent example is the estimation of exchange rate elasticities and pass-through. In this context, firms 'comply' with a specific exchange rate regime by choosing their invoicing currency – typically the exporter's or importer's currency. As shown by Gopinath et al. (2010), the degree of pass-through is not an immutable parameter but is conditional on this endogenous invoicing choice. If a researcher treats the exchange rate shock as a uniform treatment without accounting for the partial compliance of currency selection, the resulting estimates are reduced-form objects that lack policy invariance.

Overall, this paper calls for a reassessment of tariff-based estimates of the trade elasticity and for a more explicit treatment of policy uptake when defining the object of interest in both empirical estimation and structural evaluation of trade policy.

The rest of the paper is organised as follows. The next Section revises the literature on the trade elasticity and preference utilization rates. Section 3 presents the econometric problem and Section 4 describes the data. Section 5 shows the estimation results, while Section 6 discusses the interpretation of the elasticity in terms of structural parameters, and the implications of accounting for PURs in structural models of trade. Section 7 concludes.

2 Literature

A large share of the empirical literature estimates the trade elasticity using variation in import tariffs. Tariffs vary systematically across products, importers, and exporters, and are consistently reported by most countries. In their meta-analysis, Head and Mayer (2014) document that out of 622 structural estimates of the trade elasticity, 44% rely exclusively on tariffs, while an additional 26% use a combination of tariffs and freight costs, making tariffs the dominant source of identifying variation in the literature.[1]

Tariff-based estimates of the trade elasticity are typically obtained using data at the country-product level. Standard approaches control for aggregate demand- and supply-side confounders using exporter and importer fixed effects, sometimes interacted with products (e.g. Fontagné, Guimbard et al. 2022). Alternative identification strategies rely on ratio-type estimators that difference out country-specific factors, such as the "tetrads" approach used by Romalis (2007) and Head, Mayer and Ries (2010), or related methods employed by Parro (2013) and Caliendo and Parro (2015). Despite methodological differences, these

[1]Statistics computed using the online replication package of Head and Mayer (2014).

approaches share a common feature: identification comes from importers applying different tariff rates to different exporters for the same product.

This source of variation arises primarily from Preferential Trade Agreements, which allow countries to apply tariffs below the MFN rate mandated by the World Trade Organization. As a result, tariff-based estimates of the trade elasticity implicitly rely on preferential tariffs as the key identifying variation.

Despite the similarity of methods, estimated tariff elasticities display substantial heterogeneity across studies. Head and Mayer (2014) report a median elasticity of 5 for estimates based on tariffs and freight costs, with a standard deviation of 9.3. While part of this dispersion reflects sectoral heterogeneity, significant variation remains even when focusing on aggregate elasticities. As discussed in Section 3, part of this heterogeneity might be due to differences in the average preference utilization rate across samples. By ignoring partial uptake of tariff preferences, existing studies recover a eligible-tariff elasticity rather than the elasticity of trade with respect to the tariffs effectively paid.

While most of the papers reviewed by Head and Mayer (2014) use cross-sectional variation in tariffs, more recent work estimates trade elasticities using changes in tariffs over time, often within a difference-in-differences framework. This literature typically finds smaller elasticities, in the range of 1.5 to 2 (e.g. Amiti, Redding et al. (2019), Fajgelbaum et al. (2020) and Boehm et al. (2023)). These small diff-in-diff elasticities might be partly due to staggered tariff changes across countries and products. Importantly, these approaches also abstract from partial utilization of tariff preferences.

Note that the issue of non-compliance also applies when looking at variation in MFN tariffs while using PTA countries as the control group, as done in Boehm et al. (2023). In this case, we have full compliance in the treatment group (MFN countries), but some in the control group are affected as well (PTA countries trading under MFN). As for PTA tariff changes, partial compliance leads to downward biases in the estimation.

Estimated trade elasticities are widely used as key inputs in quantitative trade models. For example, Caliendo and Parro (2015) use tariff-based elasticities to evaluate the general equilibrium effects of the North America Free Trade Agreement. The elasticities of Caliendo and Parro (2015) have also been used by other authors such as Costinot and Rodríguez-Clare (2014) or Dhingra et al. (2017). As a result, issue in elasticity estimates has direct implications for counterfactual analyses.

Partial compliance extends beyond tariff-based approaches. Using the exchange rate to estimate the trade elasticity, the response of trade or prices will depend on the invoicing currency of the transaction. If an exchange rate shock is treatment, compliance is the share of trade using the subject currency. After finding marked differences between exchange rate- and tariff-based trade elasticity, Fontagné, Martin et al. (2018) says that 'the international elasticity puzzle is worse than previously thought'. However, comparing intent-to-treat effects might be misleading, as differences in compliance rates between tariffs and exchange rates will drive differences in the estimated elasticities.

Gopinath et al. (2010) provide evidence on this. The authors study how the exchange rate pass-through into US prices depends on the currency of invoice. They find that differences across countries in aggregate pass-through (which do not condition on the invoicing currency) tend to disappear once they account for the invoicing currency, as different countries have different shares of trade invoiced in US dollars. Boz et al. (2019) also provides similar evidence, showing how differences in exchange rate pass through across country-pairs can be explained by the invoicing currency shares, and Amiti, Itskhoki et al. (2022) show that the exchange rate pass-through estimated with firm-level data depends on the currency of invoice.

To the best of my knowledge, no paper estimating the trade elasticity from tariff data takes into account preference utilization rates. A likely reason is data availability: comprehensive information on imports by tariff regime is publicly available only for a limited set of countries, notably the EU and the United States. Nonetheless, given the availability of such data, it is surprising that existing studies have not incorporated preference utilization into the estimation of tariff elasticities.

At the same time, a separate literature has studied the determinants of preference utilization in depth. This work documents that utilization rates increase with the tariff margin between MFN and preferential rates (e.g. Hakobyan 2015; Legge and Lukaszuk 2024), reflecting the incentives to incur the costs of claiming preferences. Evidence suggests that utilization involves primarily fixed costs rather than variable costs, as utilization rates tend to be higher for larger trade flows (Keck and Lendle 2012; Nilsson and Dotter 2011; Hayakawa et al. 2014), and decline with more restrictive rules of origin (Cadot et al. 2006; Hakobyan 2015). Learning effects further reduce these costs over time (Krishna et al. 2025).

This paper connects these two previously separate strands of the literature. While the trade elasticity literature relies heavily on tariff variation generated by PTAs, it abstracts from the fact that these tariffs are only partially applied in practice. By explicitly accounting for preference utilization, this paper reconciles tariff-based elasticity estimates with the institutional reality of trade agreements and provides a unified framework for estimation and policy evaluation.

3 Estimating the trade elasticity with PTA tariffs

This section describes how the trade elasticity can be identified when tariff preferences are only partially utilized. I interpret preferential tariff reductions as a continuous treatment subject to one-sided non-compliance. Firms located in PTA partner countries are eligible for preferential tariffs and therefore belong to the treatment group, but only a subset of these firms actually claim the tariff preference. Firms exporting from MFN partner countries form the control group and face full compliance, as they cannot access tariff preferences.

The standard tariff-based approach used in the literature to estimate the trade elasticity relates bilateral imports to tariff rates. In its simplest form, the estimating equation can be written as:

$$\ln M_{ij} = \alpha_i + \alpha_j + \beta \ln(1 + \text{tariff}_{ij}) + \boldsymbol{\gamma} \boldsymbol{X}_{ij} + \epsilon_{ij} \tag{1}$$

where M_{ij} are imports of country j from i, α_i and α_j are exporter and importer fixed effects, and $\boldsymbol{X}_{ij}$ contains bilateral controls. Tariffs are modelled as one plus the tariff rate (e.g., a tariff of 10% is modelled as 1.1). Under standard trade models with constant elasticity, the coefficient β is interpreted as the trade elasticity and is commonly used as a structural parameter in counterfactual and welfare analysis. Importantly, in (1) the tariff is measured as the MFN tariff if the countries do not have a trade agreement, and the preferential rate if they have.

Because MFN tariffs are constant across exporters for a given importer, identification of β relies on tariff variation generated by Preferential Trade Agreements, which allow importers to apply different tariffs across exporters. Implicitly, this approach treats preferential tariffs as if they were fully applied to all eligible trade flows.

However, when trade agreements are only partially utilized, the preferential tariff assigned by the agreement differs from the tariff actually paid by exporters. In this case, regressions of imports on preferential tariffs recover the effect of tariff eligibility rather than the effect of the effective tariff. In the language of the treatment effects literature, the standard tariff regression identifies an intent-to-treat effect.

With partial utilization, preferential tariff assignment affects trade outcomes only through the subset of firms that actually claim the preference. This is a one-sided non-compliance setting: firms in MFN countries never receive treatment, while firms in PTA countries may or may not use the preferential tariff.

In aggregated data at the country or country-product level, compliance is naturally summarized by the preference utilization rate, defined as the share of imports entering under the preferential regime. The PUR therefore measures the fraction of imports that actually receives the treatment. The effective tariff change faced by exporters at the product-level is given by the tariff reduction multiplied by the PUR.

To recover the elasticity of imports with respect to applied trade costs, I adopt an instrumental variables approach. The change in preferential tariffs induced by a PTA provides variation in tariff assignment, while the interaction of the tariff change with the PUR captures the treatment effectively received.

Focusing on changes from the period before to the period after the entry into force of a PTA, the reduced-form eligible-tariff regression is given by:

$$\Delta \ln M_{ij} = \alpha_i + \alpha_j + \beta^{\text{Eligible}} \Delta \ln(1 + \text{tariff}_{ij}) \tag{2}$$

This regression captures the effect of tariff eligibility. Under standard assumptions and a constant-elasticity trade structure, the eligibility coefficient satisfies:

$$\beta^{\text{Eligible}} = \varepsilon \times \overline{PUR}$$

where ε is the structural elasticity of imports with respect to tariffs and $\overline{PUR}$ is the average preference utilization rate in the sample. As a result, ignoring partial utilization leads to an attenuation bias and makes estimates sensitive to the in-sample average PUR.

To recover ε, I estimate the following IV specification:

$$\Delta \ln M_{ij} = \alpha_i + \alpha_j + \beta^{\text{Effective}} \left[\text{PUR}_{ij} \times \Delta \ln(1 + \text{tariff}_{ij})\right] \tag{3}$$

instrumenting the interaction term with the tariff change itself:

$$\text{PUR}_{ij} \times \Delta \ln(1 + \text{tariff}_{ij}) = \alpha_i + \alpha_j + \pi \Delta \ln(1 + \text{tariff}_{ij}) \tag{4}$$

This IV estimator recovers the local average treatment effect of tariff reductions on trade flows – that is, the elasticity of imports with respect to the tariffs effectively applied.

The identifying restriction here is that, conditional on the fixed effects, tariff preferences affect imports only to the extent that they are utilized. This exclusion restriction appears quite natural. As Ornelas (2016) put it, '... preferences can promote exports only if firms use them'.

The validity of this IV strategy rests on three conditions. First, instrument relevance requires that tariff changes generate sufficient variation in the effective tariff reduction (PUR $\times$ tariff change). This is ensured mechanically: firms can only utilize preferences that are offered, and larger tariff reductions create stronger incentives for utilization.

Second, the exclusion restriction requires that tariff eligibility affects trade flows only through actual utilization. While 'news effects' or reduced uncertainty could theoretically provide alternative channels, the primary mechanism in structural gravity models operates through changes in the effective price of traded goods, which depends on utilization.

Third, the instrument must be exogenous to unobserved determinants of import growth. While a country's own PTA tariff schedules may reflect political-economy considerations – with more competitive or sensitive sectors receiving slower liberalization – I address this concern by instrumenting the effective tariff change using statutory tariff changes from other PTAs negotiated by different country pairs. These 'external' tariff schedules are plausibly uncorrelated with the focal country's sector-specific import shocks while remaining strongly correlated with its own tariff changes through the common structure of PTA liberalization.

Finally, monotonicity is a standard assumption in the LATE framework that ensures the IV estimate captures the treatment effect for compliers rather than an average of opposing behaviours. In this context, it requires that preferential tariff eligibility does not induce firms to avoid utilization – a condition naturally satisfied since claiming preferences is a cost-minimizing choice when tariff savings exceed compliance costs.

The logic above can also be derived from a heterogeneous-firm model of trade with fixed costs of accessing preferential tariffs, as in Demidova and Krishna (2008). In such a setting, only a subset of firms use the preferential regime, while others continue exporting under MFN tariffs. Total imports are therefore the sum of MFN and preferential imports. This implies that the response of aggregate imports to a tariff reduction depends both on the elasticity of firm-level sales and on the share of trade flows that actually benefit from the preferential tariff. The latter object is directly observed in the data as the preference utilization rate (PUR).

Under standard assumptions in the international trade literature – CES preferences, unbounded productivity, and ordered export cutoffs – the log-derivative of imports with respect to the tariff margin equals the structural trade elasticity ε multiplied by the PUR. As a result, regressions that relate imports to preferential tariffs but ignore utilization recover an elasticity attenuated by the average PUR in the sample.

Appendix A.1 formally derives this result and shows how the IV estimator used in this paper recovers the structural elasticity of imports with respect to tariffs.

3.1 The empirical model

This section presents the empirical specification used to estimate the trade elasticity accounting for the partial utilization of tariff preferences. The estimating equation explicitly incorporates preference utilization rates (PURs) as a measure of compliance. It is also consistent with a theoretical framework where heterogeneous firms self-select into the preferential tariff regime as presented in the Appendix.

The estimation is done on imports of the European Union by partner and product, including both imports from Free Trade Agreement (FTA) partners (the treatment group) and MFN countries, representing the control group (see Section 4 for more details on the dataset).

I estimate the dynamic response of imports to tariff liberalization using local projections. This approach is well suited to the EU trade agreements considered here, as tariff reductions are phased in over time and preference utilization increases gradually following entry into force. The local projections approach has two key advantages in this setting. First, it allows the trade elasticity to vary flexibly over time without imposing a parametric structure on the adjustment path. This is important because import responses may differ in the short and long run. Second, it naturally accommodates the staggered and phased-in nature of FTA tariff reductions, allowing each horizon h to capture the cumulative effect of liberalization from the entry into force. The baseline regression is:

$$\Delta^h \ln M_{vijt} = \alpha_{it}^h + \alpha_{jt}^h + \alpha_{vt}^h + \beta^h \left(PUR_{vij,t+h} \times \Delta^h \ln \tau_{vijt} \right) + u_{vijt}^h \tag{5}$$

where $\Delta^h x_t = x_{t+h} - x_{t-1}$. The dependent variable in equation (5) is the log-change in imports between $t-1$ and $t+h$ for a given product v, exporter i and importer j. The fixed effects absorb exporter-time, importer-time, and product-time shocks.

The key regressor is the interaction between the tariff change and the preference utilization rate at horizon $t+h$. For MFN countries, tariff changes are zero by construction. For FTA partners, tariff changes are measured relative to MFN tariffs in the period prior to entry into force, so that $\Delta^h \ln \tau_{vijt} = \ln \left(\tau_{vijt+h}^{PRF} / \tau_{vjt-1}^{MFN} \right)$ represents preferential-MFN tariff ratio. This term also represents the log of the tariff margin at time $t+h$.

Note that the PUR for the pre-treatment period is not defined, as the FTA was not in place. When estimating (5) in the pre-treatment horizons, I use the PUR of the first FTA period.

Estimating (5) presents three econometric challenges. First, the PUR is mechanically endogenous, as it is constructed from trade flows at horizon $t+h$. Second, preferential tariff schedules may be endogenous to exporter-product characteristics. Third, tariff liberalization is staggered across agreements and products. I address these issues jointly using an instrumental-variables local projections framework.

The interaction term $PUR_{vij,t+h} \times \Delta^h \ln \tau_{vijt}$, which measures the actual treatment received, is instrumented with the initial tariff change $\Delta^0 \ln \tau_{vijt}$, representing treatment assignment. This mirrors the standard LATE setup: treatment assignment (tariff eligibility) instruments actual treatment (applied tariff reductions scaled by compliance). The use of the initial change in tariffs is done in line with the local projections literature on the estimation of cumulative multipliers (Jordà and Taylor 2025). Intuitively, instrumenting the h-horizon change in the policy with its initial change allows to account for the dynamics of the policy roll-out, which in our case is given by the staged tariff reductions over time.

A remaining concern is that preferential tariff reductions may themselves be endogenous to exporter-product characteristics, for instance if the EU liberalizes more slowly in sectors where a given partner is particularly competitive. To address this type of endogeneity, I instrument the initial tariff change for country i and product v with the average initial tariff change granted by the EU to other FTA partners in the same product, adopting a leave-one-out strategy. The instrument is defined as:

$$\Delta^0 \ln \tau_{vijt}^{IV} = n^{-1} \sum_{i' \neq i} \Delta^0 \ln \tau_{vi'jt} \tag{6}$$

where n^{-1} is the number of FTAs minus one. With a little stretch to notation, the time index t on the right-hand-side of (6) indicates the time of entry into force of the FTA, not a real time. These initial tariff changes have a strong predictive power for those of the four FTA countries considered in the estimation sample, and they remove the possible exporter-specific endogeneity of preferential rates.

The identifying assumption is that tariff liberalizations reflect common EU negotiating templates and sector-level liberalization schedules that are largely harmonized across FTAs, generating cross-agreement correlation in tariff reductions. At the same time, other countries' tariff preferences are not linked to the competitiveness of a given exporter. Conditional on the fixed effects, tariff changes granted to other FTA partners should not directly affect imports from country i, except through their predictive power for the tariff schedule applied to i.

Finally, tariff liberalization is staggered across agreements and products, both because FTAs enter into force at different dates and because tariff reductions are phased in within agreements. To address this issue, I adopt a local projections difference-in-differences design in which, at each horizon h, treated observations are units whose preferential tariff reduction begins at time t, while the control group consists of units that have not yet experienced any tariff reduction by $t+h$. Units already treated prior to t are excluded from the comparison set. This ensures that identification relies exclusively on comparisons between newly treated and not-yet-treated units, avoiding contamination from already treated observations.

The approach follows recent advances in difference-in-differences estimation with staggered treatment timing (Callaway et al. 2024; Sun and Abraham 2021; Dube et al. 2025). The key insight is that valid comparisons require contrasting newly treated units against clean controls that have not yet been exposed to treatment. By excluding already-treated observations and focusing on the timing of treatment initiation, the estimator recovers an interpretable average treatment effect at each horizon that is not contaminated by treatment effect heterogeneity across cohorts.

The local projections framework integrates naturally with this design. At each horizon h, the regression estimates the cumulative effect of tariff liberalization h periods after entry into force, allowing the trade elasticity to evolve flexibly over time. This is particularly valuable given that both tariff reductions and preference utilization rates adjust gradually: firms may take time to learn about preferences, adjust their production to meet rules of origin, or renegotiate contracts. The horizon-specific estimates β^h therefore capture both the direct price effect of tariffs and the dynamic adjustment of trade flows.

This local projections difference-in-differences design is inspired by Dube et al. (2025). However, I extend their approach by allowing treatment intensity to vary over time to account for the gradual phase-in of tariff preferences. The estimator therefore combines the LP approach to estimate cumulative multipliers (Jordà and Taylor 2025, section 4) together with the staggered diff-in-diff design of Dube et al. (2025).

To summarise, the empirical approach involves the estimation of (5) instrumenting the actual treatment $PUR_{vij,t+h} \times \Delta^h \ln \tau_{vijt}$ with the average initial tariff change granted by the EU to other FTA partners, and subject to the following sample-selection to deal with staggered adoption:

$$\begin{cases} \Delta \ln \tau_{vij,t-k} = 0 \,\text{for}\, k \geq -h & \text{control} \\ \Delta \ln \tau_{vij,t} \neq 0 \,\text{and}\, \Delta \ln \tau_{vij,t-k} = 0 \,\text{for}\, k \geq 1 & \text{treatment} \end{cases} \tag{7}$$

Appendix B provides full details on the IV strategy, the sample construction, and the treatment timing restrictions.

For comparison, I also estimate the classic eligible-tariff regression that ignores preference utilization:

$$\Delta^h \ln M_{vijt} = \alpha_{it}^h + \alpha_{jt}^h + \alpha_{vt}^h + \beta_{\text{Eligible}}^h \Delta^h \ln \tau_{vijt} + u_{vijt}^h \tag{8}$$

estimated using the same IV strategy and sample selection. As shown before, this specification recovers an intent-to-treat elasticity equal to the structural elasticity multiplied by the average PUR. From the estimation of (5) and (8) we expect to find:

$$\beta^h = \frac{\beta_{\text{Eligible}}^h}{\overline{PUR}_h}$$

The next sections presents the data used and the empirical results.

4 Data

This section describes the data sources, sample construction, and key features of tariff reductions and preference utilization rates used in the empirical analysis.

Data sources: The imports data by tariff regime are obtained from Eurostat Comext at the CN 8-digit level. I retain only "normal" imports, excluding inward and outward processing trade as these flows are subject to tariff rebates, and construct total imports and preference utilization rates (PURs) for each exporter-importer-product-quarter observation.

Tariff data are from the TARIC database and aggregated from the 10- to the 8-digit level following Nilsson (2011). I extract MFN and preferential tariffs as well as special duties (anti-dumping, countervailing, safeguards). Start and end dates allow me to reconstruct the staging of FTA tariff reductions over time.

Selecting the sample: The sample covers quarterly EU imports at the CN 8-digit level from 2009–2024 for the 26 EU member states (excluding Croatia and the UK due to changes in membership).

I focus on FTAs that introduced a discrete transition from MFN to preferential tariffs. Most EU FTAs are excluded because partner countries already enjoyed preferential access prior to the agreement (e.g., GSP schemes). I further exclude agreements entering into force before 2002, because data on imports by tariff regime are not available before 2002. This leaves four FTA partners: South Korea, Canada, Japan, and Singapore.

As a control group, I include six MFN exporters (Australia, China, Hong Kong, New Zealand, Taiwan, and the United States) which face stable MFN tariffs throughout the sample. Although China was formally eligible for GSP until 2015, its utilization rate after 2003 is below 1%, and it is therefore treated as MFN.

Product classifications change over time. I construct a time-consistent product panel following Boehm et al. (2023), retaining only product codes with a one-to-one correspondence over time. Products that split or merge exit the sample at the time of reclassification. Approximately 77% of CN codes in 2009 have a one-to-one match through 2024.

For each FTA, I further restrict the sample to products that existed prior to the agreement's entry into force, ensuring that tariff changes can be clearly attributed to the FTA.

Finally, I exclude (i) products subject to special duties, (ii) products with MFN tariff changes, and (iii) products with seasonal tariffs. These restrictions ensure that identification relies exclusively on FTA-induced tariff variation.

Data description: The four FTAs entered into force at different point in times, and granted similar but not identical tariff preferences. The agreement with Korea entered into force provisionally in July 2011 (and fully in 2015). The EU-Canada FTA started in September 2017 and the one with Japan in February 2019, while the Singapore agreement entered into force in December 2019.

Table 1 reports the summary statistics of the variables used in the estimation. However, given the presence of dynamics in the tariff changes and the PUR, the figures below provide a more informative picture of the data.

Table 1: Summary statistics

variable	count	mean	std	min	p25	median	p75	max
ln(imports) FTA	5346684	8.364	3.031	0.000	6.155	8.278	10.487	20.830
ln(imports) MFN	13113515	9.154	3.113	0.000	6.903	9.214	11.383	22.732
Sub-sample with tariff preferences								
PUR	2194112	0.251	0.399	0.000	0.000	0.000	0.521	1.000
PRF tariff	7873845	0.193	1.208	0.000	0.000	0.000	0.000	62.400
MFN tariff	7873845	5.621	3.946	0.700	2.700	4.500	7.500	74.900
Sub-sample of initial tariff changes								
$\Delta \ln(1 + tariff)$	376794	-0.036	0.034	-0.559	-0.052	-0.027	-0.012	-0.000
$\Delta \ln(1 + tariff)$ IV	376794	-0.036	0.037	-0.559	-0.058	-0.027	0.000	0.000
PUR	98074	0.114	0.283	0.000	0.000	0.000	0.000	1.000
PUR $\times \Delta \ln(1 + tariff)$	98074	-0.003	0.012	-0.199	-0.000	-0.000	-0.000	-0.000

Figure 2a shows the average tariff rate since the FTA entry into force. Prior to the FTA, the average tariff was about 4% for all countries. In terms of tariff reductions, Singapore stands out as its tariff reduction was phased in more slowly. Canada had the largest immediate reduction, liberalizing about 98% of the tariff lines immediately (note: the figure is computed on products for which there are positive imports at any time during 2009-24, hence excluding products which are never imported).

Figure 2b shows the evolution of average PURs following FTA entry into force. Utilization increases gradually and stabilizes after roughly three years, but remains low on average – below 40% for Korea and below 5% for Singapore.[2]

[2]Trade-weighted PURs are substantially higher, reaching over 80% for Korea and above 60% for Japan. This gap reflects the concentration of preference use among larger exporters.

Figure 2: Tariffs and Preference Utilization Rates Over Time

(a) Average tariff over time (b) Average PUR over time

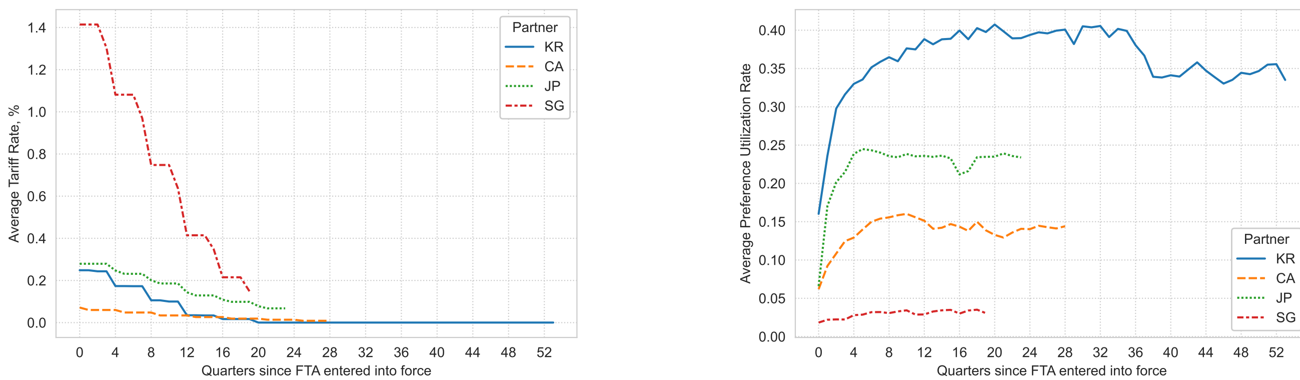

Source: author's elaboration of TARIC and Eurostat data. The preference utilization rates are computed excluding products with a zero MFN tariff or no preferential tariff margin.

The dynamics of the PUR over time will reflect two mechanisms. First, the staging of tariff reductions naturally suggests that, as more tariff preferences become available, more firms will be able and willing to use them. Second, even for immediate and one-off tariff reductions, PUR can take time to adjust due to learning (Krishna et al. (2025)) and possibly to knowledge diffusion about the FTA. It might also reflect the time taken to adjust supply chains in order to meet the rules of origin.

To assess the dynamic relationship between tariff liberalization and preference utilization, I estimate local projections of PURs on tariff reductions. This is done with a version of (12) where the PUR as the dependent variable, and restricting the sample to FTA partners only. Figure 3 shows that PURs respond strongly in the first two years following liberalization and then stabilize, consistent with both tariff staging and adjustment frictions.

Figure 3: Response of PUR to tariff liberalization

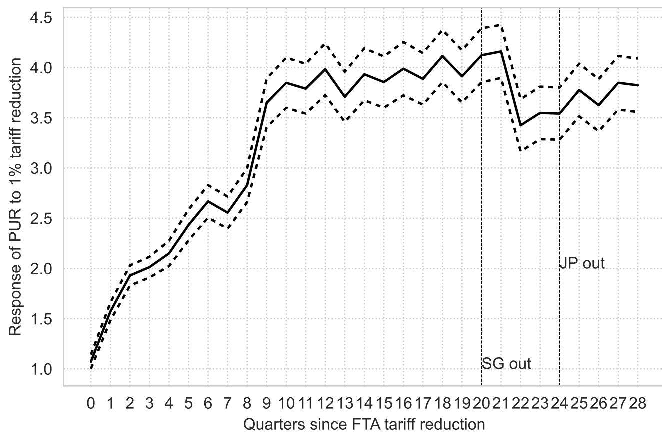

The figure reports the coefficients of a local projection regression of preference utilization rates on tariff changes, together with the 95% confidence interval. The estimation is done on the sample of FTA partner countries.

5 Results

This section presents the econometric results. I first consider a binary treatment indicator, which equals one in case of tariff liberalization, and then move to the estimation of the trade elasticity.

5.1 Effects of tariff liberalization on liberalized products

I start by estimating the average effect of tariff liberalization using a binary treatment indicator, which brings us closer to the standard micro-econometric settings. Let D_{ivt} be a dummy that equals one if product v from country i faced an FTA tariff reduction in period t and zero otherwise. I estimate a Local Projections Difference-in-Differences regression (Dube et al. 2025):

$$\Delta^h \ln M_{ijvt} = \gamma^h_{\text{Eligible}} \Delta D_{ivt} + \gamma \Delta \ln M_{ijvt-1} + \alpha^h_{it} + \alpha^h_{jt} + \alpha^h_{vt} + \epsilon^h_{ijvt} \tag{9}$$

where $\Delta^h \ln M_{ijvt}$ is the log-change in imports between $t-1$ and $t+h$, and the fixed effects absorb exporter-time, importer-time, and product-time shocks. This regression captures the effect of being eligible for preferential tariffs, which is an intent-to-treat effect, not the effect of effective tariff usage.

To account for partial utilization of the preferential tariff, I define the applied treatment as the interaction of the dummy with the preference utilization rate (PUR), $\Delta D_{ivt} \times PUR_{iv,t+h}$, and instrument it with ΔD_{ivt}:

$$\Delta^h \ln M_{ijvt} = \gamma^h_{\text{Effective}} \left(\Delta D_{ivt} \times PUR_{iv,t+h}\right) + \gamma \Delta \ln M_{ijvt-1} + \alpha^h_{it} + \alpha^h_{jt} + \alpha^h_{vt} + \epsilon^h_{ijvt} \tag{10}$$

Aggregated at the product level, compliance is measured by PUR rather than a binary indicator, so this specification recovers a weighted Local Average Treatment Effect, with more weight on larger firms. Standard LATE results imply $\gamma^h_{\text{Effective}} \approx \gamma^h_{\text{Eligible}} / \overline{PUR}_h$.

Figure 4, illustrates the results. The eligibility model ignoring PUR reports smaller effects, while the effective-tariff model accounting for PUR shows larger coefficients, consistent with the theoretical relation above. Figure 4c plots the ITT coefficients divided by the average PUR at each horizon, demonstrating near-perfect overlap with the effective-tariff estimates.

Figure 4: Effects of tariff liberalization on liberalized products

(a) Eligibility model (b) Effective-tariff model (c) Comparison of models

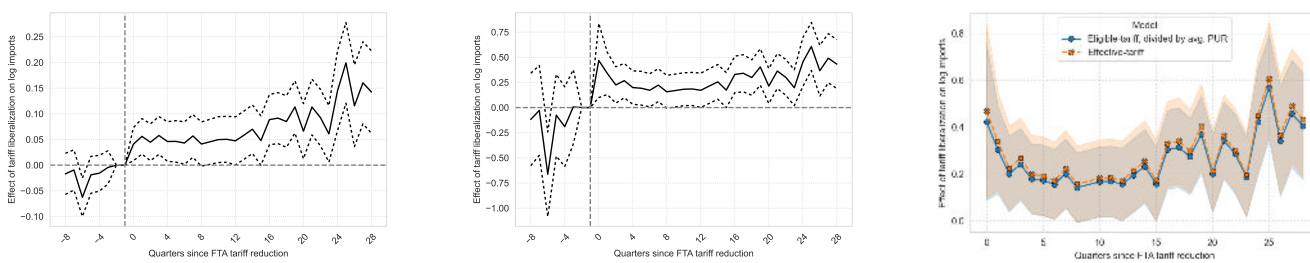

The figures report the local projection difference-in-differences estimates for the effects of tariff liberalization, together with their 95% confidence intervals. The dependent variable is horizon-change in the log of imports, and tariff liberalizations are identified with a dummy that takes value of one if a country-product saw an FTA tariff change, and zero otherwise. Figure 4a reports the results for the eligible-tariff regressions of model (9), which do not account for the preference utilization rate. The results in Figure 4b are based on model (10), where the tariff liberalization dummy is interacted with the preference utilization rate and instrumented with the tariff liberalization dummy. Figure 4c compares the two models, reporting the coefficients of Figure 4a divided by the average preference utilization rate at each horizon, and the results of Figure 4b.

5.2 The tariff elasticity

Figure 5 reports the results of the estimation of the tariff elasticity. In Figure 5a, the elasticity is estimated using the tariff as the independent variable, without accounting for the preference utilization rate. Figure 5b uses the interaction of the tariff change with the PUR, as in model (5). In both cases, the independent variable is instrumented with the average initial tariff change of other EU FTAs, and regressions include the one-period lagged change in imports.

The eligible-tariff model yields elasticities consistent with the literature: roughly -1 at horizon zero, increasing over time and reaching -3.5 seven years after FTA entry into force. In contrast, the effective-tariff model produces substantially larger elasticities: around -6 in the first two years, reaching -12.6 by horizon 28.

First-stage regressions are highly predictive, with F-statistics exceeding 1,000 across models and horizons. Detailed first-stage coefficients are reported in the Appendix.

According to the discussion presented in Section 3, the coefficient $\beta^h_{\text{Eligible}}$ of the model not accounting for PURs represents an eligible-tariff elasticity, equal to the structural trade elasticity ε multiplied by the average PUR. The coefficient β^h of the effective-tariff model recovers the structural elasticity, $\beta^h_{\text{Eligible}}/\overline{PUR}_h \approx \beta^h$.

Figure 5c illustrates this relationship. The effective-tariff elasticity closely tracks the eligible-tariff coefficient divided by the average PUR at each horizon, confirming the theoretical expectation. Minor deviations appear at the latest horizons.

These results confirm that, when tariff variation arises from FTAs and PURs are below 100%, regressions of imports on tariffs without accounting for utilization recover only an intention-to-liberalize effect. Accounting for PUR allows recovery of the structural trade elasticity, in line with the LATE interpretation of partial compliance.

Figure 5: Estimates of the dynamic tariff elasticity

(a) Eligible-tariff model (b) Effective-tariff model (c) Comparison of models

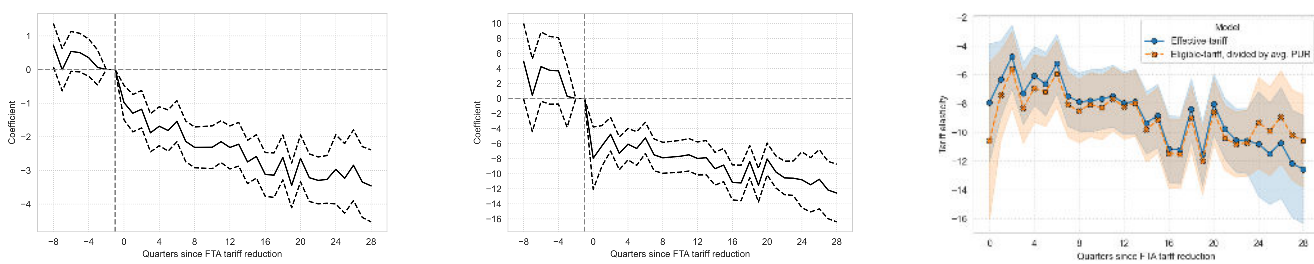

The figures report the local projection difference-in-differences estimates for the tariff elasticity, together with their 95% confidence intervals. The dependent variable is horizon-change in the log of imports. Figure 5a reports the results for the eligible-tariff elasticity of model (8), which does not account for the preference utilization rate. The results in Figure 5b are based on model (5), where the tariff change is interacted with the preference utilization rate and instrumented with the tariff change itself. Figure 5c compares the two models, reporting the coefficients of Figure 4a divided by the average preference utilization rate at each horizon, and the results of Figure 5b.

5.3 Additional results and robustness tests

I perform three robustness exercises to assess the sensitivity of the estimated tariff elasticity:

Alternative lags of the dependent variable: The baseline specification includes one lag of the change in log imports. I re-estimate the model (5) controlling for three and seven lags. The results (Figure 13) are very similar to the baseline. The specification with seven lags yields slightly more stable elasticities for horizons 20-28, stabilizing around 10.

Alternative IV strategy: The baseline uses a leave-one-out instrument based on the initial tariff changes of other EU FTAs. Omitting this instrument produces slightly smaller elasticities, although the overall dynamic pattern remains largely unchanged (Figure 13).

Heterogeneity by tariff size: The structural model (Appendix A.1) suggests that larger tariff reductions may generate larger elasticities due to a non-constant extensive margin elasticity. To explore this, I split products into above- and below-median tariff changes and re-estimate the effective-tariff regressions. Results are shown in Figure 6. Elasticities for small tariff cuts are noisier, reflecting lower variation in both the tariff change and the PUR. In contrast, larger tariff cuts yield more stable estimates, closely aligned with the baseline. Across the 29 horizons, 62% of estimates for large tariff cuts exceed those for smaller cuts in absolute value, providing suggestive evidence for the hypothesis that larger tariff reductions generate larger elasticities.

Figure 6: Heterogeneous elasticity by tariff margin

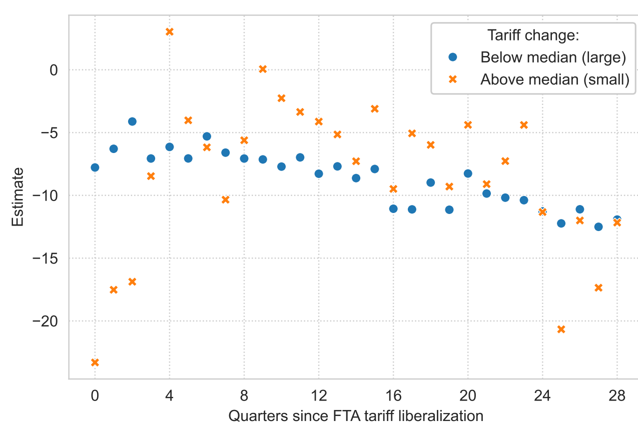

Overall, the robustness checks confirm the baseline finding, as well as the intuition that a regression of imports on statutory tariffs – without accounting for the partial uptake of preferential tariffs – recovers an elasticity with respect to eligible tariff changes rather than the effective tariff actually faced by firms.

In contrast, using an IV framework in line with the LATE literature, we can recover the effective-tariff elasticity, which is equal to the eligible-tariff elasticity divided by the average PUR of the estimation sample.

The next section discusses the mapping of econometric estimates to structural models of trade, arguing that the effective-tariff elasticity rather than the eligible-tariff one is the relevant object for trade models.

6 Implications for quantitative modelling

The previous sections established that tariff-based estimates identify different objects depending on whether preference utilization is accounted for. This section explores the implications of this distinction for quantitative trade models.

One of the main motivations for estimating tariff elasticities is their use as inputs in quantitative trade models. It is therefore essential to clarify what the eligible-tariff and effective-tariff elasticities represent, and which of the two corresponds to a structural, policy-invariant parameter. Ultimately, the goal is to identify model primitives that are stable across policy environments.

The eligible-tariff elasticity is a reduced-form object. It combines the structural trade elasticity ε with the average preference utilization rate in the estimation sample:

$$\beta^{\text{Elig}} = \varepsilon \times \overline{PUR}$$

Because this elasticity depends mechanically on the in-sample PUR, it is not portable across policy scenarios. For example, an eligibility elasticity estimated in a setting with an average PUR of 50% is not informative for an MFN tariff liberalization or a customs union, where compliance is effectively 100%. As a result, the eligible-tariff elasticity does not satisfy the Lucas critique (Lucas Jr, 1976): it varies with the policy regime and cannot serve as a primitive parameter in structural models.

By contrast, the effective-tariff elasticity estimated with the IV strategy identifies a structural elasticity ε. It captures the true response of trade flows to effectively applied tariffs, and – under the standard assumption of a constant trade elasticity – it is the appropriate parameter for quantitative modelling. While heterogeneous treatment effects imply that the effective-tariff elasticity is 'local' to firms actually using FTA tariffs, it nonetheless corresponds to the relevant primitive of models with trade-regime heterogeneity. Crucially, this elasticity is stable across policy environments: accepting the existence of a fixed ε, the effective-tariff elasticity remains the same whether estimated in samples with average PURs of 10%, 50%, or 100%.

The modelling implication is straightforward. Partial uptake of tariff preferences should enter through the effective tariff change, not through the elasticity. Standard practice implicitly, and perhaps unwillingly, scales down the trade elasticity. Because import responses are non-linear in tariffs in standard trade models, rescaling the elasticity leads overprediction of the effects of tariff liberalization. Theory-consistent modelling requires adjusting the tariff reduction itself by the relevant preference utilization rate, while keeping the elasticity as a policy-invariant structural parameter.

The remainder of the section illustrates these points in the context of the EU-Japan trade agreement. I first relate the estimated elasticity to standard trade models and discuss how preference utilization can be incorporated. I then present a quantitative exercise using two modern general equilibrium models: one that (incorrectly) scales the elasticity to capture partial uptake, and one that (correctly) scales the tariff change itself.

6.1 Interpreting the elasticity

Costinot and Rodríguez-Clare (2014) show that many micro-foundations for trade models deliver similar general-equilibrium counterfactuals. In models with perfect competition, or imperfect competition with homogeneous firms, the key parameter is the trade elasticity σ, which governs the intensive-margin response of trade flows to trade costs.

In models with firm-level heterogeneity à la Melitz (2003), two margins matter: the intensive margin, still governed by σ, and the extensive margin, governed by the parameter θ. In these settings, the elasticity of imports with respect to tariffs differs from the elasticity with respect to iceberg trade costs (e.g., transport costs), because tariff revenues are collected by the importing country.

Since the partial uptake of tariff preferences reflects firm-level heterogeneity, evaluating FTAs in structural models requires understanding both margins. It is therefore useful to have a sense of the magnitudes of σ and θ.

Under the standard Melitz-Pareto model, the total tariff elasticity can be expressed as $\sigma\theta/(\sigma - 1)$, where θ represents the Pareto shape parameter of the firm productivity distribution. In practice, we want to have an idea of the sizes of σ and θ. Costinot and Rodríguez-Clare (2014) work with an assumption that $\theta = 1.65 \times \sigma$. Crozet and Koenig (2010), using firm-level data, estimate a relation $\theta = 1.4 \times \sigma$ on aggregate trade flows.

Here, I use the dynamics of preference utilization together with the estimated horizon-specific tariff elasticities to infer the intensive- and extensive-margin elasticities. Conceptually, the total elasticity at horizon h can be written as the sum of the intensive and extensive margin contributions:

$$\text{Total elasticity}_h \approx \sigma + \text{Ext. margin response from PUR dynamics}_h$$

To back out values of the intensive and extensive margin, I set up a non-linear system relating the estimated horizon effective-tariff elasticities to the response of the PUR to a tariff reduction as shown in Figure 3. The intuition is that extensive margin response can be summarised by a function of the response of the PUR to the tariff change. In Appendix B.2, I derive an expression for the total elasticity as a function of the PUR, its response to a tariff change, the tariff margin and σ. I then ask what value of σ is consistent with the data.

I base this analysis on the effective-tariff elasticity specification controlling for seven lags of the dependent variable, which are stable in the last horizons around a value of 10. The non-linear estimation yields a central value of $\sigma \approx 5.85$. This estimate is robust to alternative specifications, including different lag structures in the local projections and the inclusion of a linear trend.

Under the standard assumption of Pareto-distributed firm productivity, we can express the total tariff elasticity – which is our estimate of the effective-tariff elasticity – as $\varepsilon = \sigma\theta/(\sigma - 1)$. Targeting $\varepsilon = 10$, a value of $\sigma = 5.85$ implies a Pareto shape parameter $\theta \approx 8.3$. Note that the ratio of these two parameters is 1.5, right in between the 1.65 considered by Costinot and Rodríguez-Clare (2014) and the 1.4 estimated by Crozet and Koenig (2010). However, the total elasticity is about twice as large as the one considered by these authors.

6.2 Quantitative evaluation of trade agreements

This section explores the implications of accounting for PUR in general equilibrium models of trade, looking at the predicted changes in trade and welfare following the entry into force of a free trade agreement.

Using the eligible-tariff elasticity has implications for policies that imply full compliance by design, such as customs unions or MFN tariff reductions. In such cases, relying on a low elasticity leads to underestimation of the effects of tariff changes. But the choice of elasticity also matters for the evaluation of FTAs where compliance is not automatic, which is the focus of this section.

For modelling FTAs, it is useful to distinguish between two approaches. The standard approach in the literature (e.g., Caliendo and Parro, 2015) can be interpreted as an eligible-tariff approach: tariff changes are measured as the full reduction from the MFN to the preferential tariff, and the elasticity used is the low eligibility elasticity, implicitly scaled by the average PUR. By contrast, the effective-tariff approach – consistent with a model where firms self-select into the preferential regime – uses the larger effective-tariff elasticity but scales the tariff change itself by the PUR. In short, the two approaches differ in where the PUR enters: the eligible-tariff method scales the elasticity; the effective-tariff method scales the tariff change.

To make the comparison between the two modelling strategies as close as possible, I adopt a Melitz-Pareto model for the eligible-tariff specification and a version of Demidova and Krishna (2008) for the effective-tariff specification. Both models feature multiple sectors, no intermediate inputs, and labour L_i is the single factor of production earning wage w_i. In both settings, the structural tariff elasticity is given by $\varepsilon = \theta\sigma/(\sigma - 1)$, where σ is the intensive-margin elasticity and θ is the Pareto shape parameter of the productivity distribution.

The two models share all structural features except for tariff-regime heterogeneity. This regime heterogeneity arises as accessing tariff preferences involves an additional fixed cost compared to the standard MFN regime, but grants lower variable costs through the preferential tariff rates. The fixed/variable costs trade-off implies that only the most productive firms access the preferential tariff regime. Importantly for the comparison, the effective-tariff model collapses to the eligible-tariff one under the restriction of no cost to access preferential tariffs. In this sense, the effective-tariff model nests the standard one.

I first review the prediction of the models in terms of partial effects – i.e., holding income and price indexes constant – and then move to the general equilibrium setting.

6.2.1 Partial effects

Before turning to the general equilibrium counterfactuals, it is useful to compare the partial equilibrium predictions of the two modelling approaches. I rely on exact hat algebra, comparing the FTA scenario to a baseline. Hats denote proportional changes, $\hat{x} = x'/x$; thus $\hat{x} = 1$ indicates no change relative to the baseline.

Let $T = \tau^{MFN}/\tau^{PRF}$ denote the margin between the MFN and preferential tariff, with $\tau^{MFN} = 1 + \text{MFN tariff}$ and $\tau^{PRF} = 1 + \text{PRF tariff}$.

The standard eligible-tariff model predicts predicts the partial equilibrium import response as:

$$\hat{M}^{\text{Eligible}} = T^{\varepsilon^{\text{Eligible}}} = T^{\varepsilon \times \overline{PUR}} \tag{11}$$

since this method applies the full tariff change but uses the scaled elasticity $\varepsilon^{\text{Eligible}} = \varepsilon \times \overline{PUR}$.

In contrast, the effective-tariff model predicts:

$$\hat{M}^{\text{Effective}} = 1 + (T^{\sigma} - 1)\, s \tag{12}$$

where s is the (endogenous) size-share of firms that use the preferential regime. While s is not directly observable, it can be recovered from the observed PUR given σ and the tariff margin as shown in (22) in the Appendix. Importantly, we have $s = 0$ when $PUR = 0$ and $s = 1$ when $PUR = 1$.[3]

[3] The measure s can be interpreted as the export share of firms using the FTA if they were to face the same tariff as the

One could argue that the partial uptake of tariff preferences is accounted for in the scaled elasticity, as $\varepsilon^{\text{Eligible}} = \varepsilon \times PUR$. However, because the import response is non-linear in tariffs, scaling the elasticity does not replicate the mechanics of the effective-tariff model. As a result, the eligible-tariff model systematically predicts larger changes in imports than the effective-tariff model.

A closed-form analytical comparison of (11) and (12) is not possible, so I conduct two numerical exercises. Throughout, I fix $\varepsilon = 10$.

I first fix the tariff margin and the utilization rate close to their sample averages (respectively 1.05 and 30%), and consider different values of σ from 2 to 6. The eligibility model response, which does not depend on σ, predicts a change in imports of 15.8%. The effective-tariff model instead predicts changes between 2.8% and 9.5%, depending on σ.

Second, I experiment with different values of the tariff margin and PUR, fixing $\sigma = 5.85$. In this set of evaluation, I consider the PUR and tariff margin structurally related, such that a higher tariff margin produces a higher PUR. Guided by the results on the response of the PUR to a tariff reduction presented in Figure (3), I assume a relation $PUR = 4 \times \ln T$. Letting T vary between 1.02-1.25, the eligibility model always overpredicts the partial trade effects compared to the effective-tariff model. For the largest tariff margin of 1.25, the eligibility model over-predicts compared by the effective-tariff model by a factor of 3.4. The full set of results is reported in Table 3a in the Appendix.

6.2.2 General equilibrium

For the comparison of the general equilibrium effects, I provide numerical results for the EU-Japan trade agreement, as both parties report data on PUR for the agreement.

The GE analysis is based on a static multi-sector model without intermediate inputs, with firm heterogeneity and labour as the only factor of production. Consumption across different sectors is aggregated with a Cobb-Douglas function. Details of the modelling are reported in the Appendix Section E.

The standard eligible-tariff model is the multi-sector Melitz-Pareto model with tariffs reviewed in Costinot and Rodríguez-Clare (2014). Although the effective-tariff model does not yield a gravity equation which is log-linear in tariffs, the general equilibrium counterfactuals can be computed with exact hat algebra as in the class of models reviewed by Costinot and Rodríguez-Clare (2014).

For the eligible-tariff model, I solve the following system of non-linear equations:

firms not using the FTA.

$$\hat{Y}_i Y_i = \sum_k \sum_j \lambda_{ijk} \frac{\hat{\tau}_{ijk}^{1-\varepsilon^{\text{Eligible}}} \left(\hat{Y}_i\right)^{-\theta+\frac{-\theta+\sigma-1}{\sigma-1}}}{\sum_{i'=1}^{n} \lambda_{i'jk} \hat{\tau}_{i'jk}^{1-\varepsilon^{\text{Eligible}}} \left(\hat{Y}_{i'}\right)^{-\theta+\frac{-\theta+\sigma-1}{\sigma-1}}} \beta_{jk} \frac{Y_j \hat{Y}_j}{1-\pi'_j} \tag{13}$$

where $\lambda_{ijk} = X_{ijk}/E_{jk}$ is the expenditure share of country j sector k on exports from country i, β_{jk} is the expenditure share of country j in sector k, and π_j represents tariff revenues in country j. The tariff change $\hat{\tau}_{ijk}$ is equal to the MFN/PRF tariff margin for EU-Japan trade ($\hat{\tau}_{ijk} = \tau_{ijk}^{PRF}/\tau_{ijk}^{MFN}$) and $\hat{\tau}_{ijk} = 1$ for all other countries. Note that the exponent on the tariff change is $1 - \varepsilon^{\text{Eligible}}$ as the counterfactual modelling is done on tariff-inclusive trade flows.

I consider two approaches for measuring the tariff elasticity $\varepsilon^{\text{Eligible}}$. In the first case, labelled 'Eligibility no-PUR' I consider $\varepsilon^{\text{Eligible}}$ as a constant, and fix $\varepsilon^{\text{Eligible}} = \theta\sigma/(\sigma-1) = 3.5$ to target the horizon-28 estimates from the eligible-tariff regressions reported in Figure 5a. This is the standard modelling approach.

In the second case, labelled 'Eligibility+PUR' I follow the same logic used for the comparison of the partial effects, and set $\varepsilon_{ijk}^{\text{Eligible}} = PUR_{ijk} \times \varepsilon$, fixing the structural parameter to $\varepsilon = 10$. Here I let the elasticity vary by country-pair, so to reflect the different PURs of EU exports to Japan and EU imports from Japan. This exercise allows me to incorporate the information on the PUR in the eligible-tariff model, by scaling the tariff elasticity in relation to the preference utilization rate.

For the effective-tariff model, the system of non-linear equations to solve is:

$$\hat{Y}_i Y_i = \sum_k \sum_j \lambda_{ijk} \frac{\left(\hat{Y}_i\right)^{-\theta-\frac{\theta-\sigma+1}{\sigma-1}} \left[1 + \left(T_{ijk}^{\sigma\left[\iota_{ijk}\frac{\theta}{\sigma-1}+(1-\iota_{ijk})\right]-1} - 1\right) s_{ijk}\right]}{\sum_{i'} \lambda_{i'jk} \left(\hat{Y}_{i'}\right)^{-\theta-\frac{\theta-\sigma+1}{\sigma-1}} \left[1 + \left(T_{i'jk}^{\sigma\left[\iota_{i'jk}\frac{\theta}{\sigma-1}+(1-\iota_{i'jk})\right]-1} - 1\right) s_{i'jk}\right]} \beta_{jk} \frac{Y_j \hat{Y}_j}{1-\pi'_j} \tag{14}$$

where $T_{ijk} = \tau_{ijk}^{MFN}/\tau_{ijk}^{PRF}$ is the tariff margin for EU and Japan trade, while it takes value of one for all other countries. The PUR enters the equation in the calculation of the shares s_{ijk} which are computed as in (22) in the Appendix. The term ι is a dummy that equals one if the PUR is 100%, as in such case the effective-tariff model collapses to the standard one – hence the tariff elasticity is $1-\sigma\theta/(\sigma-1)$ as $s_{ijk} = 1$.

Both (13) and (14) are a system of N equations in N unknowns (national income Y_i). As one of the equation is redundant, I set one country as the numeraire with $\hat{Y}_i = 1$. Prior to running the FTA counterfactual exercise, I use the model equations to remove aggregate trade imbalances as explained in Costinot and Rodríguez-Clare (2014).

In terms of the model parameters, for the effective-tariff model I target the structural parameter

$\varepsilon = \frac{\sigma\theta}{\sigma-1} = 10$ with the latest horizons elasticity of the seven-lag specification. Based on the results of Section 6.1, I set $\sigma = 5.85$ implying $\theta = 8.3$. The same values are used for the eligible-tariff model with the tariff elasticity depending on the PUR.

However, the same set of values for σ and θ do not square well with setting $\varepsilon^{\text{Eligible}} = \sigma\theta/(\sigma-1) = 3.5$ in the 'Eligibility no-PUR' version. To maintain some consistency with the other modelling designs, I keep $\theta/(\sigma-1)$ constant at 1.71 (=8.3/(5.85-1)). We then obtain $\sigma = 2.05$ and $\theta = 1.79$. Obviously, the choice of different parameter values for σ and θ makes the comparison of the two models more difficult to interpret, and can have a potentially big impact on the welfare effects. For this reason, the most interesting comparison of between the effective-tariff and the Eligibility+PUR models – the latter accounting for the PUR in the tariff elasticity. Table 2 summarises the parameter choices.

Table 2: General equilibrium model parameters

Parameter	Model		
	Eligibility no-PUR	Eligibility+PUR	Effective-tariff
σ	2.05	5.85	5.85
θ	1.79	8.3	8.3
tariff elasticity	$\frac{\sigma\theta}{\sigma-1} = 3.5$	$\frac{\sigma\theta}{\sigma-1} \times PUR_{ijk}$	-
PUR	-	data for EU-Japan	data for EU-Japan

The counterfactual exercise for the EU-Japan trade agreement is based on trade and production data from the FIGARO tables for 2018. The data have been aggregated to 34 sectors and 13 countries. I aggregated the EU27 to a single country so that the comparison across models is more straightforward. The tariff data are taken from WITS.

The results of the simulations for the three modelling approaches are reported in Figure 7. Each point represents the predicted percentage change in exports from EU to Japan (Figure 7a) and from Japan to the EU (Figure 7b). The blue circles represent the results for the effective-tariff model. The red triangles those for the Eligibility+PUR model, with the elasticity varying with the PUR. The crosses stands for the Eligibility no-PUR model with a constant tariff elasticity. For some sectors, the Eligibility+PUR model overpredicts the trade effects by a factor of 2. Variation across sectors in the overpredictions is due to sectors having different PURs.

Apart for sectors with no changes in tariffs, for which the trade effects are only the results of GE forces (e.g., Japan does not have tariffs on motor vehicles C29), the effective-tariff model predicts smaller effects of the tariff liberalization on exports. Aggregating across all sectors, the Eligibility+PUR model predicts a 1% more exports from the EU to Japan and 4% more for Japan to the EU than the effective-tariff model

Next, a central question to trade models is the impact of trade liberalizations on welfare. In this class of models, welfare can be expressed as real consumption. Changes in welfare can be computed with the modified ACR formula (Arkolakis et al. (2012)) to account for multiple sectors and tariff revenues:

$$\hat{C}_j = \frac{1-\pi_j}{1-\pi_j'} \prod_k \left(\hat{\lambda}_{jj,k}\right)^{-\beta_{jk}/\theta} \tag{15}$$

where $\frac{1-\pi_j'}{1-\pi_j}$ measures the change in tariff revenues and $\hat{\lambda}_{jj,k}$ is the change in domestic consumption shares, aggregate across sectors with the sectoral expenditure shares β_{jk}. With tariff revenues, large values of θ plays an even more important role in downsizing the welfare effects than in the standard ACR formula, as welfare also depends on tariff revenues.

My simulations for a tariff-only FTA between the EU and Japan show very little welfare changes, which are negative for both the effective-tariff and the Eligibility+PUR models, while it is positive for the Eligibility no-PUR model with a constant tariff elasticity. Anyway, in all cases the welfare changes are not distinguishable from 1 if we look at the first four decimal places. The reasons for these small welfare effects are due to the tariff liberalizations removing small initial tariffs (the pre-FTA average across sectors is around 2%), and the fact that tariff revenues decrease. Indeed, while domestic consumption shares fall for both the EU and Japan, pushing towards a welfare increase, they are counterbalanced by reductions in tariff revenues.

Figure 7: General equilibrium trade effects of the EU-Japan FTA, tariffs only

(a) EU exports to Japan (b) Japan exports to the EU

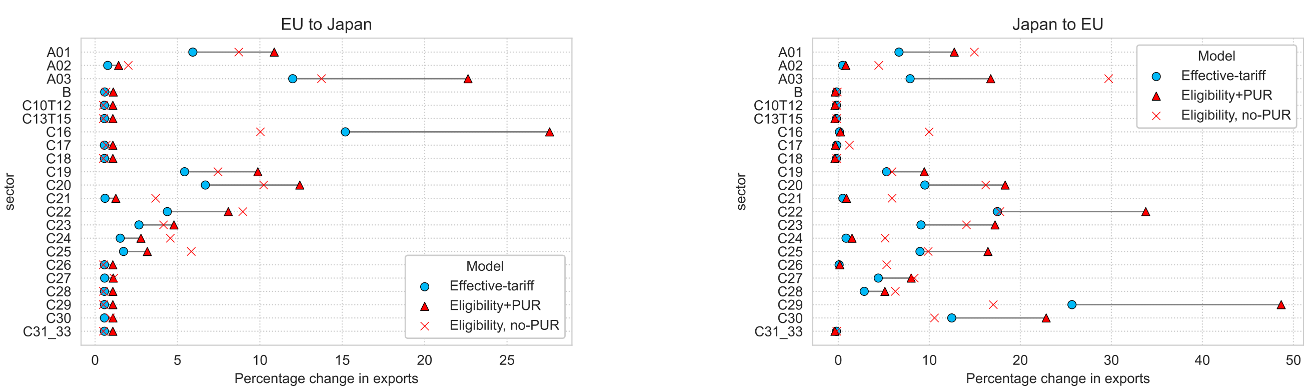

The figure reports the predicted general equilibrium percentage changes in exports between the EU and Japan following the FTA tariff liberalization. Other trade costs are held fixed. The blue circles represent the results for the effective-tariff model. The red triangles those for the eligible-tariff models with the tariff elasticity depending on the PUR, while the crosses are for the eligibility model with a constant tariff elasticity.

Finally, it is worth noting that for ex-ante evaluations of FTAs, preference utilization rates are not known. While this is certainly true, the same could be argued for changes in trade costs induced by FTAs, which are often assumed or somehow estimated. The same can be done with PURs. For instance, analysts could use machine learning tools, which excel in making out-of-sample predictions, to estimate expected PURs based on the available data from other agreements given the negotiated tariffs and other product and country characteristics.

7 Conclusions

The paper reconciles the structural modelling, tariff elasticity and PUR literatures. Recognizing that when PUR is less than 100% we are in situation of non-compliance with treatment assignment, we can borrow the tools developed for the estimation of the Local Average Treatment Effect literature. In particular, we obtain that the standard regression of imports on tariffs exploiting variation in FTA tariff preferences to estimate the trade elasticity can only recover an intent-to-treat effect, or an eligible-tariff elasticity, which differ from the structural trade elasticity required for counterfactual modelling. Moreover, as the eligible-tariff elasticity depends on the average PUR of the estimation sample, it can have little external validity when applied to different settings than the estimation sample.

Instead, an instrumental variable regression where the effective tariff change, given by the tariff change multiplied by the preference utilization rate, is instrumented with the tariff change itself (representing treatment assignment) we can recover the structural policy-invariant trade elasticity. The eligible-tariff method, which yields an elasticity of 3.5, is biased towards zero by partial compliance. Addressing the issues yield elasticities between 10 and 12.

The paper also shows the implications of not accounting for PUR in the structural modelling of free trade agreements. Not accounting for PURs generally leads to an overestimation of the trade effect of a tariff liberalization, depending on the value of the PUR. Taking the EU-Japan agreement to a multi-sector general equilibrium model, I find that the standard modelling approach can overpredict the trade effects by a factor of 2.

While this paper focuses on tariff variation from PTAs and the role of preference utilization, in a companion study exploiting the UK's MFN tariff reform of 2021 (*Tamberi, forthcoming CITP Working Paper*), I find similarly large elasticities. Unlike this paper, the companion paper exploits pre-reform exposure to MFN tariffs of PTA countries and mechanical changes in the UK's MFN tariff schedule to deal with the potential endogeneity of the MFN tariff changes. The convergence in the trade elasticity values across two independent policy settings reinforces the external validity of my results and suggests that the true trade elasticity is substantially larger than previously thought.

The paper calls for a reassessment of the available estimates of the tariff elasticity. It also calls for modellers of free trade agreements to make a more consistent use of tariff preference utilization. As data on preference utilization rates are not readily available for all countries, effort should be put in creating a global PUR dataset.

References

Amiti, Mary, Oleg Itskhoki and Jozef Konings (2022). "Dominant currencies: How firms choose currency invoicing and why it matters". In: *The Quarterly Journal of Economics* 137.3, pp. 1435–1493.

Amiti, Mary, Stephen J Redding and David E Weinstein (2019). "The impact of the 2018 tariffs on prices and welfare". In: *Journal of Economic Perspectives* 33.4, pp. 187–210.

Arkolakis, Costas, Arnaud Costinot and Andrés Rodríguez-Clare (2012). "New trade models, same old gains?" In: *American Economic Review* 102.1, pp. 94–130.

Boehm, Christoph E, Andrei A Levchenko and Nitya Pandalai-Nayar (2023). "The long and short (run) of trade elasticities". In: *American Economic Review* 113.4, pp. 861–905.

Boz, Emine, Gita Gopinath and Mikkel Plagborg-Møller (May 2019). "Dollar Invoicing and the Heterogeneity of Exchange Rate Pass-Through". In: *AEA Papers and Proceedings* 109, pp. 527–32.

Cadot, Olivier et al. (2006). "Product-specific rules of origin in EU and US preferential trading arrangements: an assessment". In: *World Trade Review* 5.2, pp. 199–224.

Caliendo, Lorenzo and Fernando Parro (2015). "Estimates of the Trade and Welfare Effects of NAFTA". In: *The Review of Economic Studies* 82.1, pp. 1–44.

Callaway, Brantly, Andrew Goodman-Bacon and Pedro HC Sant'Anna (2024). *Difference-in-differences with a continuous treatment*. Tech. rep. National Bureau of Economic Research.

Costinot, Arnaud and Andrés Rodríguez-Clare (2014). "Trade theory with numbers: Quantifying the consequences of globalization". In: *Handbook of International Economics*. Vol. 4. Elsevier, pp. 197–261.

Crozet, Matthieu and Pamina Koenig (2010). "Structural gravity equations with intensive and extensive margins". In: *Canadian Journal of Economics/Revue canadienne d'économique* 43.1, pp. 41–62.

Demidova, Svetlana and Kala Krishna (2008). "Firm heterogeneity and firm behavior with conditional policies". In: *Economics Letters* 98.2, pp. 122–128.

Dhingra, Swati et al. (2017). "The costs and benefits of leaving the EU: trade effects". In: *Economic Policy* 32.92, pp. 651–705.

Dube, Arindrajit et al. (2025). "A Local Projections Approach to Difference-in-Differences". In: *Journal of Applied Econometrics*.

Fajgelbaum, Pablo D et al. (2020). “The return to protectionism”. In: *The Quarterly Journal of Economics* 135.1, pp. 1–55.

Felbermayr, Gabriel J, Benjamin Jung and Mario Larch (2013). “Icebergs versus tariffs: a quantitative perspective on the gains from trade”. In.

Fontagné, Lionel, Houssein Guimbard and Gianluca Orefice (2022). “Tariff-based product-level trade elasticities”. In: *Journal of International Economics* 137, p. 103593.

Fontagné, Lionel, Philippe Martin and Gianluca Orefice (2018). “The international elasticity puzzle is worse than you think”. In: *Journal of International Economics* 115, pp. 115–129.

Gopinath, Gita, Oleg Itskhoki and Roberto Rigobon (2010). “Currency choice and exchange rate pass-through”. In: *American Economic Review* 100.1, pp. 304–336.

Hakobyan, Shushanik (2015). “Accounting for underutilization of trade preference programs: The US Generalized System of Preferences”. In: *Canadian Journal of Economics/Revue canadienne d’économique* 48.2, pp. 408–436.

Hayakawa, Kazunobu, Hansung Kim and Hyun-Hoon Lee (2014). “Determinants on utilization of the Korea–ASEAN free trade agreement: margin effect, scale effect, and ROO effect”. In: *World Trade Review* 13.3, pp. 499–515.

Head, Keith and Thierry Mayer (2014). “Gravity equations: Workhorse, toolkit, and cookbook”. In: *Handbook of International Economics*. Vol. 4. Elsevier, pp. 131–195.

Head, Keith, Thierry Mayer and John Ries (2010). “The erosion of colonial trade linkages after independence”. In: *Journal of International Economics* 81.1, pp. 1–14.

Jordà, Òscar and Alan M Taylor (2025). “Local projections”. In: *Journal of Economic Literature* 63.1, pp. 59–110.

Keck, Alexander and Andreas Lendle (2012). “New evidence on preference utilization”. In: *World Trade Organization Staff Working Paper No. ERSD-2012-12*.

Krishna, Kala et al. (2025). “Learning to use trade agreements”. In: *Journal of Development Economics* 174, p. 103424.

Legge, Stefan and Piotr Lukaszuk (2024). “The firm-level costs of utilizing free trade agreements”. In: *International Economics* 178, p. 100484.

Lucas Jr, Robert E (1976). “Econometric policy evaluation: A critique”. In: *Carnegie-Rochester conference series on public policy*. Vol. 1. North-Holland, pp. 19–46.

Melitz, Marc J (2003). “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity”. In: *Econometrica* 71.6, pp. 1695–1725.

Nilsson, Lars (2011). “Principles of EU Imports, Tariffs, and Tariff Regimes”. In: *Journal of World Trade* 45.4.

Nilsson, Lars and Caroline Dotter (2011). "Small trade flows, compliance costs and trade preferences: The case of EU imports from African LDCs". In: *Economics Discussion Paper* 2011-42.

Ornelas, Emanuel (2016). "Special and differential treatment for developing countries". In: *Handbook of commercial policy*. Vol. 1. Elsevier, pp. 369–432.

Parro, Fernando (2013). "Capital-skill complementarity and the skill premium in a quantitative model of trade". In: *American Economic Journal: Macroeconomics* 5.2, pp. 72–117.

Romalis, John (2007). "NAFTA's and CUSFTA's Impact on International Trade". In: *The Review of Economics and Statistics* 89.3, pp. 416–435.

Sun, Liyang and Sarah Abraham (2021). "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects". In: *Journal of Econometrics* 225.2, pp. 175–199.

Appendix

A Theory

A.1 Representation of the elasticity in a structural model of trade

If we consider a situation in which not all firms use the trade agreement, we necessarily have firm heterogeneity. Demidova and Krishna (2008) provide a tractable model of trade with firm-heterogeneity *a la* Melitz (2003). As accessing tariff preferences involves an additional fixed costs, only the most productive firms use the trade agreement. At the same time, we could think of a model where firms, either because they make different sourcing decisions or because of they differ in their technologies, they might be able to meet the rules of origins and hence access preferences.

Here I provide some insights without imposing too many restrictions on the supply-side of model. In order to derive a tractable model, we have to proceed with the assumption that the productivity cutoffs for exporting under the MFN and preferential regime are ordered, and that the productivity distribution is unbounded. In addition, I assume CES preferences for consumers as standard in the literature. The assumptions are:

A1. CES preferences;

A2. Unbounded firm productivity distribution;

A3. The MFN and preferential productivity cutoffs for exporting are ordered.

The model presented below abstracts from trade costs other than tariffs, but these can be easily introduced and also made different for the MFN and preferential regime. None of these additions would change the main result.

Let φ index firm productivity, and define φ_{vij}^{MFN} as the cutoff above which firms export under MFN and φ_{vij}^{PRF} the one for the preferential regime. I assume that either $\varphi_{vij}^{MFN} \leq \varphi_{vij}^{PRF}$, so that firms with productivity between φ_{vij}^{MFN} and φ_{vij}^{PRF} export under MFN and those with productivity above φ_{vij}^{PRF} export under the preferential regime. Or, if $\varphi_{vij}^{MFN} \geq \varphi_{vij}^{PRF}$, firms with productivity between φ_{vij}^{PRF} and φ_{vij}^{MFN} export under the preferential regime and firms above φ_{vij}^{MFN} export under MFN.

In the following exposition, $G(\varphi)$ denotes the cumulative density function of productivity (which is unbounded) and D_j the CES aggregate demand (including income and price index). To find total imports from country i to j, we integrate over the productivity distribution. Under the assumption that the MFN and PRF cutoffs are ordered, we can derive expression (16) independently of which of the two cutoffs is the largest one.

When $\varphi_{vij}^{MFN} < \varphi_{vij}^{PRF}$, imports under the MFN regime are integrated between the two cutoffs:

$$M_{vij}^{MFN} = D_j \times N_i \left[G(\varphi_{vij}^{PRF}) - G(\varphi_{vij}^{MFN})\right] \times \left[\frac{1}{G(\varphi_{vij}^{PRF}) - G(\varphi_{vij}^{MFN})} \int_{\varphi_{vij}^{MFN}}^{\varphi_{vij}^{PRF}} \varphi^{\sigma-1} dG(\varphi)\right] \times \left(\tau_{vj}^{MFN}\right)^{-\sigma}$$

while preferential imports are those of firms with productivity above φ_{vij}^{PRF}:

$$M_{vij}^{PRF} = D_j \times N_i \left[1 - G(\varphi_{vij}^{PRF})\right] \times \left[\frac{1}{1 - G(\varphi_{vij}^{PRF})} \int_{\varphi_{vij}^{PRF}}^{\infty} \varphi^{\sigma-1} dG(\varphi)\right] \times \left(\tau_{vij}^{PRF}\right)^{-\sigma}$$

Then total imports are:

$$\begin{gathered} M_{vij} = D_j \times N_i \times \left[\tilde{\varphi}_{vij}^{MFN} \left(\tau_{vj}^{MFN}\right)^{-\sigma} + \tilde{\varphi}_{vij}^{PRF} \left(\tau_{vij}^{PRF}\right)^{-\sigma}\right] = \\ D_j \times N_i \times \tilde{\varphi}_{vij} \left(\tau_{vj}^{MFN}\right)^{-\sigma} \left[\frac{\tilde{\varphi}_{vij}^{MFN}}{\tilde{\varphi}_{vij}} + \frac{\tilde{\varphi}_{vij}^{PRF}}{\tilde{\varphi}_{vij}} \left(\tau_{vj}^{MFN}/\tau_{vij}^{PRF}\right)^{\sigma}\right] \end{gathered}$$

with $\tilde{\varphi}_{vij} = \int_{\varphi_{vij}^{min}}^{\infty} \varphi^{\sigma-1} dG(\varphi)$ representing and the average productivity over the full range of exporting firms and $\varphi_{vij}^{min} = \min[\varphi_{vij}^{MFN}, \varphi_{vij}^{PRF}]$. The ratios of $\tilde{\varphi}$ represents the size-shares of the firms in each tariff regime.

When instead we have $\varphi_{vij}^{MFN} > \varphi_{vij}^{PRF}$, MFN imports are:

$$M_{vij}^{MFN} = D_j \times N_i \left[1 - G(\varphi_{vij}^{MFN})\right] \times \left[\frac{1}{1 - G(\varphi_{vij}^{MFN})} \int_{\varphi_{vij}^{MFN}}^{\infty} \varphi^{\sigma-1} dG(\varphi)\right] \times \left(\tau_{vj}^{MFN}\right)^{-\sigma}$$

and preferential imports:

$$M_{vij}^{PRF} = D_j \times N_i \left[G(\varphi_{vij}^{PRF}) - G(\varphi_{vij}^{MFN})\right] \times \left[\frac{1}{G(\varphi_{vij}^{PRF}) - G(\varphi_{vij}^{MFN})} \int_{\varphi_{vij}^{PRF}}^{\varphi_{vij}^{MFN}} \varphi^{\sigma-1} dG(\varphi)\right] \times \left(\tau_{vij}^{PRF}\right)^{-\sigma}$$

Aggregating we obtain the same expression for total imports:

$$M_{vij} = D_j \times N_i \times \tilde{\varphi}_{vij} \left(\tau_{vj}^{MFN}\right)^{-\sigma} \left[\frac{\tilde{\varphi}_{vij}^{MFN}}{\tilde{\varphi}_{vij}} + \frac{\tilde{\varphi}_{vij}^{PRF}}{\tilde{\varphi}_{vij}} \left(\tau_{vj}^{MFN}/\tau_{vij}^{PRF}\right)^{\sigma}\right]$$

Note that when the minimum exporting cutoff is the MFN one, which is not affected by the tariff margin, then $\tilde{\varphi}_{vij}$ is not affected either.

Let $s_{vij} = \tilde{\varphi}_{vij}^{PRF}/\tilde{\varphi}_{vij}$ and $1 - s_{vij} = \tilde{\varphi}_{vij}^{MFN}/\tilde{\varphi}_{vij}$ denote the size-share of firms trading under the preferential and MFN regime, respectively. We can express total imports as the sum of MFN and preferential imports:

$$M_{vij} = M_{vij}^{MFN} + M_{vij}^{PRF} = D_j \times N_i \times \tilde{\varphi}_{vij} \left[(1 - s_{vij})(\tau_{vj}^{MFN})^{-\sigma} + s_{vij}(\tau_{vij}^{PRF})^{-\sigma}\right] = \quad (16)$$

$$D_j \times N_i \times \tilde{\varphi}_{vij}(\tau_{vj}^{MFN})^{-\sigma} \left[1 + \left(T_{vij}^{\sigma} - 1\right) s_{vij}\right] \quad (17)$$

where $T_{vij} = \tau_{vj}^{MFN}/\tau_{vij}^{PRF}$ represents the tariff margin. Note that the preference utilization rate is given by:

$$PUR_{vij} = \frac{M_{vij}^{PRF}}{M_{vij}} = \frac{T_{vij}^{\sigma} s_{vij}}{1 + \left(T_{vij}^{\sigma} - 1\right) s_{vij}} \quad (18)$$

Expression (16) tells us that imports under the preferential regime are equal to MFN imports multiplied by the factor $[1 + \left(T_{vij}^{\sigma} - 1\right) s_{vij}]$ measuring the increase in sales due to the lower FTA tariffs and their utilization. Moreover, (16) tells us that imports are not a log-linear functions of tariffs when an FTA is present.

The elasticity of imports to the tariff margin: To study the elasticity of imports to the tariff margin, I focus on the case where $s_{vij} < 1$, so that imports under both the MFN and preferential regimes are observed. Moreover, I make the additional assumption:

A4. When $s_{vij} < 1$, the exporting cutoff for the preferential regime is larger than the MFN one.

Assumption A4 implies that the minimum exporting cutoff is the MFN one, making $\tilde{\varphi}_{vij}$ independent of the tariff margin. This behaviour occurs when accessing tariff preferences involves some fixed costs as in Demidova and Krishna (2008). This would be in line with the literature on the PURs, which also finds some empirical support for the hypothesis that tariff preferences are used by larger exporters.

To see the relation between imports and tariffs, take the log of (16) and differentiate with respect to $\ln T_{vij}$:

$$\begin{aligned}\frac{\partial \ln M_{vij}}{\partial \ln T_{vij}} = \sigma \frac{T_{vij}^{\sigma} s_{vij}}{1+\left(T_{vij}^{\sigma}-1\right) s_{vij}} + \frac{\partial s_{vij}}{\partial T_{vij}} T_{vij} \frac{T_{vij}^{\sigma}-1}{1+\left(T_{vij}^{\sigma}-1\right) s_{vij}} = \\ \sigma \frac{T_{vij}^{\sigma} s_{vij}}{1+\left(T_{vij}^{\sigma}-1\right) s_{vij}} + \frac{\partial s_{vij}}{\partial T_{vij}} \frac{T_{vij}}{s_{vij}} \frac{\left(T_{vij}^{\sigma}-1\right) s_{vij}}{1+\left(T_{vij}^{\sigma}-1\right) s_{vij}} = \\ \left(\sigma + \frac{\partial s_{vij}}{\partial T_{vij}} \frac{T_{vij}}{s_{vij}}\right) \frac{T_{vij}^{\sigma} s_{vij}}{1+\left(T_{vij}^{\sigma}-1\right) s_{vij}} - \frac{\partial s_{vij}}{\partial T_{vij}} \frac{T_{vij}}{s_{vij}} \frac{T_{vij}^{\sigma} s_{vij}}{1+\left(T_{vij}^{\sigma}-1\right) s_{vij}} T_{vij}^{-\sigma} \\ = \left[\sigma + \varepsilon^{s}\left(1-T_{vij}^{-\sigma}\right)\right] \times PUR_{vij} = \bar{\varepsilon} \times PUR_{vij} \quad (19)\end{aligned}$$

where in the second line I multiplied and divided the second term by s_{vij}, in the third line I collected the common terms and multiplied and divided the last term by T_{vij}^{σ}, and in the final line I substituted the definition of the PUR in (18) and defined the elasticity of s_{vij} with respect to the tariff margin as $\frac{\partial s_{vij}}{\partial T_{vij}} \frac{T_{vij}}{s_{vij}} = \varepsilon^{s} > 0$.

With the standard Melitz-Pareto assumption, the model becomes the one of Demidova and Krishna (2008), the extensive margin elasticity is a constant $\sigma \frac{\theta - \sigma + 1}{\sigma - 1}$ (i.e., it does not depend on T^{σ}) hence $\bar{\varepsilon}$ becomes a constant equal to $\frac{\sigma\theta}{\sigma-1}$. This simplifies the expression to:

$$\frac{\partial \ln M_{vij}}{\partial \ln T_{vij}} = (\sigma + \varepsilon^{s}) \times PUR_{vij} = \varepsilon \times PUR_{vij} = \frac{\sigma\theta}{\sigma - 1} \times PUR_{vij} \quad (20)$$

Expression (20) tells us that if we regress the log of imports on the log of the tariff, what we can hope to recover from the tariff coefficient is the combination of the structural parameter ε multiplied by the average preference utilization rate of the estimation sample: $\beta^{\text{Elig}} = \varepsilon \times \overline{PUR}_{vij}$. This result is in line with the LATE literature, and shows that β^{Elig} does not represent a structural parameter.

If we consider the more general derivation in (19), then the tariff elasticity should also vary with the level of tariff liberalization. In particular, a larger liberalization (high T_{vij}) will have larger extensive margin effects magnifying the tariff elasticity. What we can recover then is an average elasticity $\bar{\varepsilon}$.

B Empirical model

B.1 Derivation of the baseline regression model

I rely on local projections to estimate the tariff elasticity in a triple-difference setting with three-way fixed effects. The approach can be motivated by looking at the change in imports following the entry into force of an FTA. Based on the theory presented in Section A.1, before the entry into force of the FTA, imports from country i can only enter under the MFN regime. Letting $t-1$ denote the pre-FTA period, we have:

$$M_{vij,t-1} = A_{vij} \times D_{jt-1} \times S_{it-1} \times C_{vt-1} \times \left(\tau_{vjt-1}^{MFN}\right)^{-\varepsilon} \times \epsilon_{vijt-1}$$

where the capital letters absorb aggregate demand and supply shocks, also at the product level, and ϵ_{vijt-1} is an error term. As the FTA enters into force at time t, we can take the ratio of imports at time $t+h$ relative to $t-1$:

$$\frac{M_{vij,t+h}}{M_{vij,t-1}} = \hat{D}_{j,t+h}\hat{S}_{i,t+h}\hat{C}_{v,t+h}\left(\hat{\tau}_{vj,t+h}^{MFN}\right)^{-\varepsilon} \times \left[1 + \left(T_{vij,t+h}^{\sigma} - 1\right) s_{vij,t+h}\right] \times \hat{\epsilon}_{vijt+h} \tag{21}$$

In the estimation, I focus only on products with a constant MFN tariff over the sample period, so that $\hat{\tau}_{vj,t+h}^{MFN}$ can be eliminated. Equation (21) is used to justify the empirical model (5) presented in the main text. Taking logs we obtain a local projections representation, with the term $\ln\left[1 + \left(T_{vij,t+h}^{\sigma} - 1\right) s_{vij,t+h}\right]$ approximated with its derivative $\varepsilon\left(PUR_{vijt+h} \times \Delta \ln \tau_{vijt+h}\right)$.

B.2 The intensive and extensive margin elasticities

This section provides the detailed derivation of how the estimated tariff elasticity can be decomposed into intensive- and extensive-margin effects, and how we recover the structural intensive-margin elasticity σ from the data.

From the structural model of Section A.1, the h-horizon tariff elasticity can be written as:

$$\beta^h = \sigma + \overline{\frac{\partial \ln s_{vij,h}}{\partial \ln T_{vij,h}}\left(1 - T_{vij,h}^{-\sigma}\right)}$$

where:

- σ is the intensive-margin elasticity of trade,

- $s_{vij,h}$ is the size share of firms using the preferential tariff at horizon h,
- $T_{vij,h} = \tau_{vj}^{MFN}/\tau_{vij}^{PRF}$ is the tariff margin
- the overline denotes the sample average over products and exporters.

The second term captures the contribution of the extensive margin, i.e., the response of selection into the preferential regime.

Inverting the definition of the preference utilization rate (18), we have:

$$s_{vij,h} = \frac{PUR_{vij,h}}{PUR_{vij,h} + (1 - PUR_{vij,h})\, T_{vij,h}^{\sigma}} \tag{22}$$

We can then express the elasticity of $s_{vij,h}$ with respect to the tariff margin in terms of the PUR:

$$\frac{\partial \ln s_{vij,h}}{\partial \ln T_{vij,h}} = \frac{\beta^{PUR,h}}{PUR_{vij,h}} - \frac{\beta^{PUR,h} + \sigma T_{vij,h}^{\sigma}(1 - PUR_{vij,h}) - \beta^{PUR,h} T_{vij,h}^{\sigma}}{PUR_{vij,h} + (1 - PUR_{vij,h})\, T_{vij,h}^{\sigma}}$$

where $\beta^{PUR,h} = \partial PUR_{vij,h}/\partial \ln T_{vij,h}$ is the coefficient from a LP regression of the PUR on the log of the tariff margin, as those presented in Figure 3.

Substituting the expression for the extensive margin elasticity into the total elasticity, we have a non-linear system:

$$\beta^{h} = \sigma + \left[\frac{\bar{\beta}^{PUR,h}}{\overline{PUR}_{vij,h}} - \frac{\bar{\beta}^{PUR,h} + \sigma \bar{T}_{vij,h}^{\sigma}\left(1 - \overline{PUR}_{vij,h}\right) - \bar{\beta}^{PUR,h} \bar{T}_{vij,h}^{\sigma}}{\overline{PUR}_{vij,h} + \left(1 - \overline{PUR}_{vij,h}\right) \bar{T}_{vij,h}^{\sigma}}\right]\left(1 - \bar{T}_{vij,h}^{-\sigma}\right) \tag{23}$$

for $h = 1, ..., H$, where $\bar{\beta}^{PUR,h}$ is the coefficient from a LP regression of the PUR on the log of the tariff margin, as those presented in Figure 3. $\overline{PUR}_{vij,h}$ is the average h-horizon PUR, and $\bar{T}_{vij,h}$ is the average h-horizon tariff margin.

Focussing on the post-FTA horizons, we have 28 estimates of β^{h} that can be related to the LHS of (23). This gives us a system of 28 non-linear equations where the only unknown is σ. I therefore apply a non-linear least square estimator to find the value of σ that minimizes the difference between the RHS and LHS of (23):

$$\min_{\sigma} \left[\beta^{h} - \sigma - \left(\frac{\bar{\beta}^{PUR,h}}{\overline{PUR}_{vij,h}} - \frac{\bar{\beta}^{PUR,h} + \sigma \bar{T}_{vij,h}^{\sigma}\left(1 - \overline{PUR}_{vij,h}\right) - \bar{\beta}^{PUR,h} \bar{T}_{vij,h}^{\sigma}}{\overline{PUR}_{vij,h} + \left(1 - \overline{PUR}_{vij,h}\right) T_{vij,h}^{\sigma}}\right)\left(1 - \bar{T}_{vij,h}^{-\sigma}\right)\right]^{2} \tag{24}$$

The results of the estimation of (24) are reported in Figure 8, plotting the absolute value of the tariff elasticity coefficients (solid black line) and the fitted values from the non-linear estimation (dashed red line). I also fit a version of (24) including a linear trend (dotted green line). For the tariff elasticity, I rely on the estimate controlling for seven lags in the log-change of imports, as it exhibits a slightly more stable behaviour.

Figure 8: Heterogeneous elasticity by tariff margin

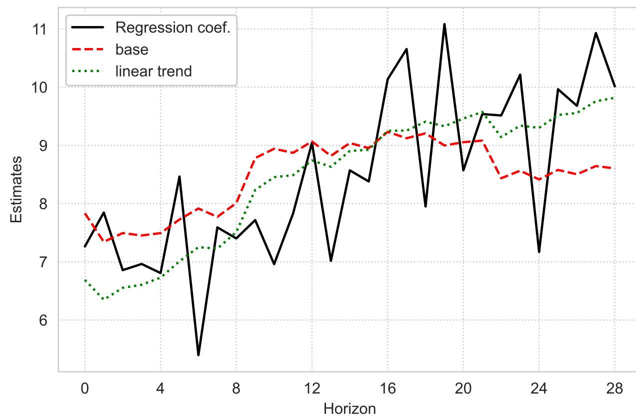

The estimation without a trend yields a value of $\sigma = 6.6$, and does a fairly decent jobs in tracking the tariff elasticity estimates. Adding a linear trend improves the fit for the latest horizons, and yields an estimate for σ of 5.3.

Using different specifications of the tariff elasticity model yields similar results. Using the baseline model with one lag, I obtain $\sigma = 6.6$ and 4.1 when adding a linear trend. The specification with three lags gives $\sigma = 7.2$ and 5.3 with a linear trend. The median across all six values if 5.85 (6.6 without linear trend and 5.3 with the trend).

If we were to take the model of Demidova and Krishna (2008), where only the most productive firms self-select into the preferential regime and the productivity distribution is Pareto, the tariff elasticity estimated by (5) is a constant equal to $\theta\sigma/(\sigma-1)$, with θ being the shape parameter of the Pareto distribution. With $\sigma = 5.85$, and a long-run elasticity of 10 we would then have $\theta = 8.29$.

C Data appendix

This section provides additional summary statistics of the the data.

The main text showed that the average and total PUR can differ substantially across countries. This differences can derive from multiple sources. Sectoral specialization will be one driver, as well as product-specific rules of origin and other factors. 9 shows the average PURs and tariff margins by product groups. A scatter plot of PUR vs tariff margins indicates a positive correlation between the tariff margin and the PUR. The tariff margin by itself tells only part of the story, as the utilization rate will also depend on how easily firms can meet rules of origin. For instance, the relatively high PUR for Minerals is likely to reflect the low complexity of these products, making it easy to meet the RoOs requirement.

To inspect the relation between PURs and rules of origin, Figure 10 plots the average PUR for Korea and Canada by type of RoOs.[4] There are quite some differences between Korea and Canada in terms of ranking. Wholly obtained rules, which affects Agrifood products, get the highest PUR in Canada but not in Korea. For Korea, the highest PUR are those with RoOs related to specific processes requirements, while rules based on the net weight of inputs into the products rank low in Korea but high in Canada.

These differences can be related to the product specialization of each country as well as the type of technology and input sourcing used by each country. For instance, maple syrup is very likely to be originating in Canada, less so in Korea. Trying to account for all these factors can be difficult as we would need some multi-dimensional index. Luckily though, all this information is condensed in one indicator: the Preference Utilization Rate.

Figure 9: PUR and tariff margins by product groups

(a) PUR by product group (b) Tariff margins by product group (c) PURs and the tariff margin

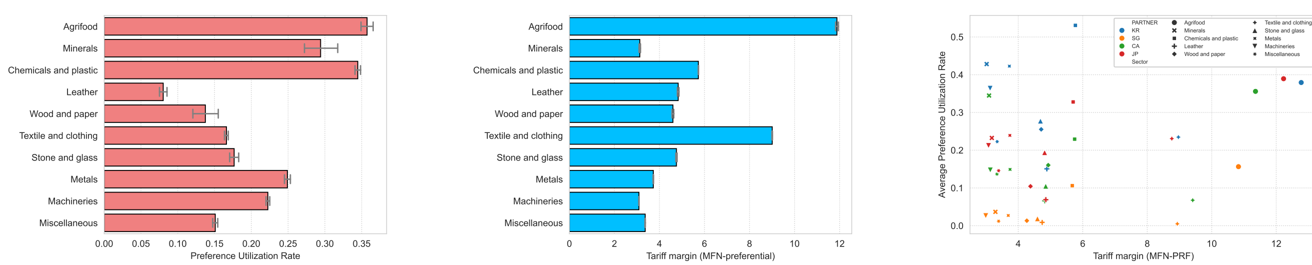

Source: author's elaboration of Eurostat and TARIC data. All data reported in the figures refer to 2024. The bar plots report the average and 95% confidence interval of the preference utilization rate and tariff margin (difference between MFN and preferential tariff). The scatter plot shows the average PUR by country and product group plotted against the average tariff margin.

[4]I am grateful to Manuel Tong for sharing the data on rules of origin.

Figure 10: Preference Utilization and Rules of Origin

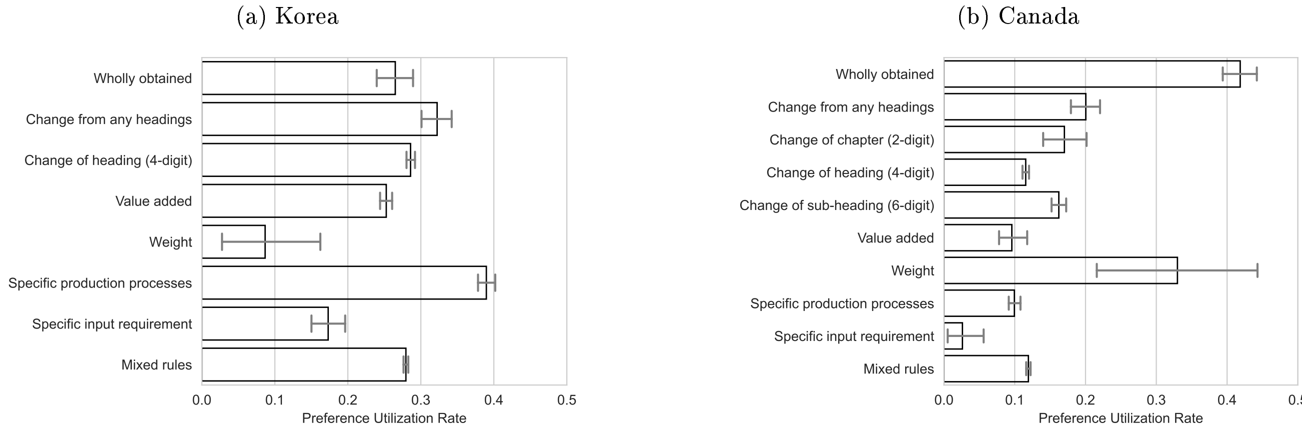

Source: author's elaboration of Eurostat and UKTPO data. The figure shows the average preference utilization rate by type of Rules of Origin, together with the 95% confidence interval.

For the product concordance over time, I follow the approach of Boehm et al. (2023). Using the correlation tables of the Combined Nomenclature (CN) across years, I map products with a 1:1 match across different vintages of the CN classification. Products that split or merge are considered as individual products. For instance, if product x in the CN 2011 classification splits into x1 and x2 in 2012, these are treated as different products. This means that x ends in 2011, and x1 and x2 start in 2021.

This approach can be conservative, but ensures that tariff changes over time do not arise because of product codes merging/splitting. The procedure will reduce the number of products observed over long time periods, hence reducing the precision of the long horizon elasticities. However, a quite large fraction of the products have 1:1 match throughout the whole sample period. Figure 11 shows the fraction of product codes declarable in any given year over time – i.e., the share of 1:1 matches. For instance, the line starting in 2010 indicates the products of the 2009 CN vintage. The big jumps occurs with the major revisions of the Harmonised System classification, which occurred in 2012, 2017 and 2022.

Figure 11: Product concordance over time

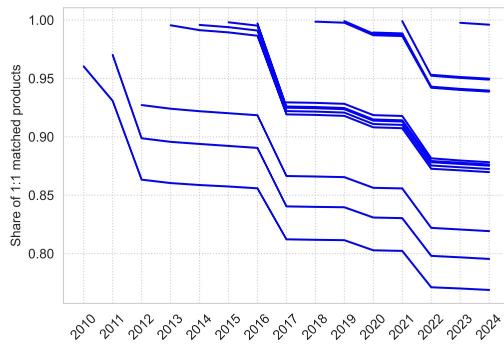

D Additional results

Figure 12: First-stage regression coefficients

(a) Model without PUR interaction

(b) Model with PUR interaction

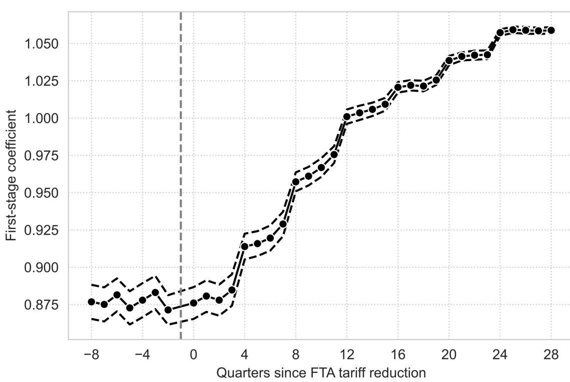

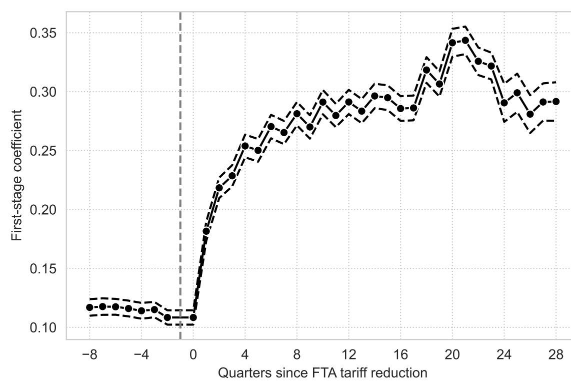

Figure 13: Robustness: IV and lag specifications

(a) Alternative specifications

(b) Comparison of models, controlling for 3 lags

(c) Comparison of models, controlling for 7 lags

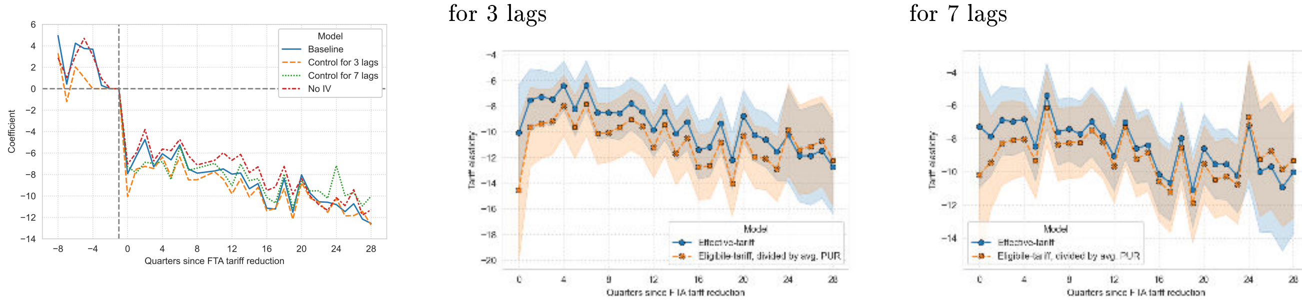

D.1 Partial equilibrium effects

This section provides the details of the calculation of partial effects of the eligibility and effective-tariff models. As a direct analytical comparison of (11) and (12) is not possible, I proceed with two numerical comparisons.

In the first comparison I hold fixed the tariff margin $T = 1.05$ and the PUR=0.3, which are close to the sample averages, and look at different values of σ. For the structural parameter ε, I target the horizons 20-28 elasticities from the regression with seven lags, setting $\varepsilon = 10$, giving $\varepsilon^{\text{Eligible}} = 10 \times 0.3 = 3$. The term s_{vij} is computed as in (22), hence it varies with σ.

In the second type of comparison I hold σ fixed, but let T vary. In this case, we should recognise that the PUR is not invariant to the tariff margin, but increases with T. Targeting the long-run estimates of the regression of PUR on $\ln T$ reported in Figure (3), I compute $PUR = 4 \times \ln T$.

For this second comparison, I let the elasticity vary with the PUR. The thought experiment is as follows. Imagine that we were to estimate the tariff elasticity from different samples, each having a different average tariff margin and PUR, with the two being related. Then I ask what would be the predicted partial effects of the average tariff liberalization in the sample. To do so, I fix the structural elasticity $\varepsilon = 10$, and compute the elasticity for the eligibility model as $\varepsilon^{\text{Eligible}} = \varepsilon \times PUR$.

The results of the partial-effect comparison are reported in Table 3. In Table 3a I let σ vary. This does not affect on the partial effect of the eligibility model, which predicts an increase in imports by 15.8%. The effects predicted by the effective-tariff model are substantially smaller, ranging from 2.9% with $\sigma = 2$ and reaching 8.2% with $\sigma = 6$.

In Table 3b I look at different tariff margins, letting the PUR adjust to the tariff margin and therefore the elasticity for the eligibility model $\varepsilon^{\text{Eligible}}$, which increases with the PUR. The eligible-tariff model predicts larger partial effects than the effective-tariff model for every values of the tariff margin. Moreover, the differences become larger as the tariff margin increases. For a tariff margin of 1.02, the eligibility model predicts an effect 1.9 larger than the effective-tariff model. When the tariff margin is 1.25, the eligibility model effects are 3.3 times larger than the effective-tariff ones. If instead we were to consider a constant elasticity for the eligibility model, setting $\varepsilon^{\text{Eligible}} = 3.5$ to match the results reported in Figure 5a, we would obtain larger partial effects for small tariff margins (+7.2% when the tariff margin is 1.02) and smaller ones when the tariff margin gets large (+118% when the tariff margin is 1.25).

Table 3: Comparison of partial effects

(a) Changing σ

σ	Partial effect on imports, % Model	
	Eligible-tariff	Effective-tariff
2	15.8	2.9
3	15.8	4.3
4	15.8	5.6
5	15.8	6.9
6	15.8	8.2

(b) Changing the tariff margin

Tariff margin	PUR	Partial effect on imports, % Model	
		Eligible-tariff	Effective-tariff
1.02	0.08	1.6	0.9
1.05	0.20	10.0	5.1
1.1	0.38	43.8	19.5
1.2	0.73	278.0	91.7
1.25	0.89	632.8	186.2

E General equilibrium modelling

This section explains the general equilibrium modelling behind the results presented in Section 6.

The model features firm heterogeneity in productivity φ. There are N countries, with the subscripts i indicating the exporter and j the importer, and multiple sectors indexed by $k = 1, ..., K$. The presentation follows closely Costinot and Rodríguez-Clare (2014). I start introducing the (standard) eligibility model, which also serves as the base for MFN trade, and then introduce the modelling of FTAs accounting for preference utilization rates. As some of the results presented are standard in the literature, I do not show all derivations step-by-step.

The eligible-tariff model: Each consumer in country j maximises consumption C_j:

$$C_j = \Pi_{k=1}^{K} C_{j,k}^{\beta_{j,k}} \tag{25}$$

where $\beta_{j,k} \geq 0$ are consumption shares and $\sum_k \beta_{jk} = 1$. The sector-specific consumption is given by:

$$C_{jk} = \left[\int_{\omega \in \Omega} c_{j,k}(\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right]^{\frac{\sigma}{\sigma-1}} \tag{26}$$

where $\sigma > 1$ is the elasticity of substitution across varieties. Each variety is sourced from one country only hence $C_{jk} = \left(\sum_i C_{ij,k}^{(\sigma-1)/\sigma}\right)^{\frac{\sigma}{\sigma-1}}$, with $C_{ijk} = \left(\int_{\omega \in \Omega_{ijk}} c_{j,k}(\omega)^{\frac{\sigma-1}{\sigma}} d\omega\right)^{\frac{\sigma}{\sigma-1}}$. The associated consumer price index is $P_j = \Pi_{k=1}^{K} P_{jk}^{\beta_{jk}}$ and the sector price indexes are given by $P_{jk} = \left(\sum_i P_{ij,k}^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$ and $P_{ijk} = \left(\int_{\omega \in \Omega_{ijk}} p_{j,k}(\omega)^{1-\sigma} d\omega\right)^{\frac{1}{1-\sigma}}$. Note that the price index is tariff-inclusive, as it represent the price faced by consumers in the importing country.

On the production side, firms are heterogeneous in their productivity φ and operate in monopolistic competition. The profit-maximising behaviour gives the optimal price as $p_{ik} = \frac{\sigma}{\sigma-1}\frac{w_i}{\varphi}$. Adding bilateral trade costs $\delta_{ijk} \geq 1$ we have the pre-tariff bilateral price given by $p_{ijk} = p_{ik}\delta_{ijk}$. When selling to market j, a firm faces a demand $c_{ijk} = p_{ijk}^{-\sigma} \beta_{jk} E_j P_{jk}^{\sigma-1}$. Here E_j is total expenditure in country j and $\beta_{jk}E_j$ is expenditure in sector k.

The firm also faces tariffs denoted by $\tau_{ijk} = 1 + \text{tariff}_{ijk}$. However, while tariff reduce demand by increasing the price consumers face, the firm does not accrue the tariff revenue, which is collected at the border. Hence firm revenues from selling to j are:

$$r_{ijk}(\varphi) = \left(\frac{p_{ijk}}{P_{jk}}\right)^{1-\sigma} \tau_{ijk}^{-\sigma} \beta_{jk} E_j = \left(\frac{\sigma}{\sigma-1} \frac{w_i \delta_{ijk}}{\varphi P_{jk}}\right)^{1-\sigma} \tau_{ijk}^{-\sigma} \beta_{jk} E_j \tag{27}$$

with profits given by $\pi_{ijk}(\varphi) = r_{ijk}(\varphi)/\sigma - w_i f_{ijk}$, where $w_i f_{ijk}$ are the fixed exporting costs paid in terms of domestic labour. Then the productivity cut-off for exporting is:

$$\varphi_{ijk}^{*} = \left(\frac{\sigma}{\sigma-1} \frac{w_i \delta_{ijk}}{P_{jk}}\right) \left(\frac{w_i f_{ijk} \sigma}{\tau_{ijk}^{-\sigma} \beta_{jk} E_j}\right)^{\frac{1}{\sigma-1}} \tag{28}$$

To get total sector-level exports of i to j we integrate firm sales over productivity. Productivity is distributed Pareto with a cumulative density function $G(\varphi) = 1 - \varphi^{-\theta}$. Total tariff inclusive exports can then be expressed with the sector-level gravity equation:

$$X_{ijk} = \frac{\tau_{ijk}^{1-\frac{\sigma\theta}{\sigma-1}} M_i \left(w_i \delta_{ijk}\right)^{-\theta} \left(w_i f_{ijk}\right)^{\frac{-\theta+\sigma-1}{\sigma-1}}}{\sum_{i'=1}^{n} \tau_{i'jk}^{1-\frac{\sigma\theta}{\sigma-1}} M_{i'} \left(w_{i'} \delta_{i'jk}\right)^{-\theta} \left(w_{i'} f_{i'jk}\right)^{\frac{-\theta+\sigma-1}{\sigma-1}}} \beta_{jk} E_j \tag{29}$$

The tariff elasticity of tariff-inclusive trade is $1 - \frac{\sigma\theta}{\sigma-1}$, where $-\frac{\sigma\theta}{\sigma-1}$ is the elasticity that can be recovered from the regression of imports on tariffs. As noted in Felbermayr et al. (2013), the mass of entrants M_i is proportional to the fixed labour endowment L_i, hence it does not change in comparative statics. Moreover, as labour is the only factor of production, total revenues are given by total income $R_i = Y_i$.

Total expenditure in country j is the sum of income $Y_j = \sum_i \sum_k X_{jik}/\tau_{jik}$, where I divide by tariffs as I defined X_{ijk} as tariff-inclusive exports, and tariff revenues $tr_j = \sum_i \sum_k \frac{\tau_{ijk}-1}{\tau_{ijk}} X_{ijk}$. When trade data are unbalanced, we must also account for the aggregate trade deficit D_j. Hence expenditure are:

$$E_j = Y_j + D_j + tr_j \tag{30}$$

Noting that tariff revenues have a recursive structure, we can re-write sectoral expenditure as:

$$E_j = \frac{Y_j + D_j}{1 - \pi_j} \tag{31}$$

where $\pi_j = \sum_k \beta_{jk} \pi_{j,k}$ represents the tariff revenues share, aggregating the sectoral shares $\pi_{j,k} = \sum_i \frac{\tau_{ij,k}-1}{\tau_{ij,k}} \lambda_{ij,k}$.

The next subsection lays down the model with trade regime heterogeneity, where firms self-select into using an FTA.

The effective-tariff model: When an FTA grants conditional tariff preferences, not all firms use them. The following model follows the approach of Demidova and Krishna (2008), which, to the best of my knowledge, is the only model which accounts for the endogeneity of PURs.

When an FTA is present, firms from country i can export to j under the MFN of preferential tariff regime. If they export under the MFN regime, the firm faces an exporting fixed cost f_{ijk} (as in the previous section) and faces the MFN tariff τ_{ijk}^{MFN}. I maintain the index i to indicate that for domestic sales ($i = j$) the tariff is zero and $\tau_{iik}^{MFN} = 1$. Alternatively, the firm can pay an additional fixed cost f_{ijk}^{PRF} to access the preferential tariff regime. If it chooses to do so, it faces the lower preferential tariff $\tau_{ijk}^{PRF} < \tau_{ijk}^{MFN}$. This setting generates a trade-off between variable and fixed costs, and two possible equilibria arise.

If the fixed costs of the preferential regime are low or the preferential tariff is substantially smaller than the MFN one, then all firms export under the preferential regime. This is a rather uninteresting case as the standard modelling of tariffs can capture all the effects. Nonetheless, the counterfactual analysis will take into account the possibility of all firms trading under the preferential regime, which happens when the PUR is 100%.

On the other hand, if the fixed costs of the preferential regime are high or the preferential tariff margin is low, only the most productive firms self-select into using the FTA. In this case, the exporting cut-off for exporting under MFN is given by (28) with the tariff being the MFN rate:

$$\varphi_{ijk}^{MFN} = \left(\frac{\sigma}{\sigma-1}\frac{w_i\delta_{ijk}}{P_{jk}}\right)\left[\frac{f_{ijk}w_i\sigma}{\beta_{jk}E_j\left(\tau_{ijk}^{MFN}\right)^{-\sigma}}\right]^{\frac{1}{\sigma-1}} \tag{32}$$

A firm choosing the MFN regime will receive the following revenues:

$$r_{ijk}^{MFN}(\varphi) = \left(\frac{\sigma}{\sigma-1}\frac{w_i\delta_{ijk}}{\varphi P_{jk}}\right)^{1-\sigma}\left(\tau_{ijk}^{MFN}\right)^{-\sigma}\beta_{jk}E_j \tag{33}$$

while the firm choosing the preferential regime receives:

$$r_{ijk}^{PRF}(\varphi) = \left(\frac{\sigma}{\sigma-1}\frac{w_i\delta_{ijk}}{\varphi P_{jk}}\right)^{1-\sigma}\left(\tau_{jk}^{PRF}\right)^{-\sigma}\beta_{jk}E_j \tag{34}$$

It is easy to add different variable trade costs δ_{ijk} for the preferential regime, but in this presentation I keep the focus on tariffs only. The two regimes also differ in fixed costs: for the MFN regime, the firm pays $f_{ijk}w_i$, while for the preferential regime it pays $(f_{ijk} + f_{ijk}^{PRF})w_i$. These additional costs can be modelled as additive or multiplicative without loss of generality.

Given the availability of the MFN tariff regime and the higher fixed costs of the preferential tariff regime, a firm uses the FTA only if:

$$r_{ijk}^{PRF}(\varphi)/\sigma - (f_{ijk} + f_{ijk}^{PRF})w_i \geq r_{ijk}^{MFN}(\varphi)/\sigma - f_{ijk}w_i \tag{35}$$

with equality indicating indifference. To find the cut-off for the preferential regime, replace the inequality with an equality in (35) and solve for productivity. This gives:

$$\varphi_{ijk}^{PRF} = \underbrace{\left(\frac{\sigma}{\sigma-1}\frac{w_i\delta_{ijk}}{P_{jk}}\right)\left[\frac{f_{ijk}w_i\sigma}{\beta_{jk}E_j\left(\tau_{ijk}^{MFN}\right)^{-\sigma}}\right]^{\frac{1}{\sigma-1}}}_{\text{MFN cut-off}}\left[\frac{f_{ijk}^{PRF}/f_{ijk}}{\left(\tau_{ijk}^{MFN}/\tau_{jk}^{PRF}\right)^{\sigma}-1}\right]^{\frac{1}{\sigma-1}} \tag{36}$$

Equation (36) shows that the cut-off for the preferential tariff regime is the MFN one shifted to the left by the term $\left[\frac{f_{ijk}^{PRF}/f_{ijk}}{\left(\tau_{ijk}^{MFN}/\tau_{jk}^{PRF}\right)^{\sigma}-1}\right]^{\frac{1}{\sigma-1}}$ which represents the trade-off between fixed and variable costs between the two regimes. This implies that, if the fixed costs of the preferential regime f_{ijk}^{PRF} are sufficiently high or if the preferential tariff margin is sufficiently low, then only the most productive firms choose the MFN regime.

To derive total imports with an FTA and heterogeneous trade regimes, we have to integrate firm MFN firms' sales from φ_{ijk}^{MFN} to φ_{ijk}^{PRF}, and preferential firms' sales from φ_{ijk}^{PRF} to infinity:

$$X_{ijk} = \underbrace{\int_{\varphi_{ijk}^{MFN}}^{\varphi_{ijk}^{PRF}} r_{ijk}^{MFN}(\varphi) dG(\varphi)}_{\text{MFN exports}} + \underbrace{\int_{\varphi_{ijk}^{PRF}}^{\infty} r_{ijk}^{PRF}(\varphi) dG(\varphi)}_{\text{Preferential exports}}$$

Similarly, the tariff-inclusive consumer price index for the sector is:

$$P_{jk} = \sum_{i=1}^{N} \left[\int_{\varphi_{ijk}^{MFN}}^{\varphi_{ijk}^{PRF}} p_{ijk}^{MFN}(\varphi, \tau^{MFN})^{1-\sigma} dG(\varphi) + \int_{\varphi_{ijk}^{PRF}}^{\infty} p_{ijk}^{PRF}(\varphi, \tau^{PRF})^{1-\sigma} dG(\varphi) \right]$$

Maintaining the assumption of Pareto distribution for productivity we obtain the following expression for tariff-inclusive MFN exports:

$$X_{ijk}^{MFN} = \left(\frac{\sigma}{\sigma - 1} \frac{w_i \delta_{ijk}}{P_{jk}} \right)^{1-\sigma} \beta_{jk} E_j M_i \frac{\theta}{\theta - \sigma + 1} \left(\tau_{ijk}^{MFN} \right)^{1-\sigma} \left(\varphi_{ijk}^{MFN} \right)^{-\theta+\sigma-1} \left[1 - \left(\frac{f_{ijk}^{PRF}/f_{ijk}}{\left(\tau_{ijk}^{MFN}/\tau_{jk}^{PRF} \right)^{\sigma} - 1} \right)^{\frac{-\theta+\sigma-1}{\sigma-1}} \right] \tag{37}$$

while tariff-inclusive preferential exports are:

$$X_{ijk}^{PRF} = \left(\frac{\sigma}{\sigma - 1} \frac{w_i \delta_{ijk}}{P_{jk}} \right)^{1-\sigma} \beta_{jk} E_j M_i \frac{\theta}{\theta - \sigma + 1} \left(\tau_{ijk}^{PRF} \right)^{1-\sigma} \left(\varphi_{ijk}^{MFN} \right)^{-\theta+\sigma-1} \left[\frac{f_{ijk}^{PRF}/f_{ijk}}{\left(\tau_{ijk}^{MFN}/\tau_{jk}^{PRF} \right)^{\sigma} - 1} \right]^{\frac{-\theta+\sigma-1}{\sigma-1}}$$

Defining $s_{ijk} = \left[\frac{f_{ijk}^{PRF}/f_{ijk}}{\left(\tau_{ijk}^{MFN}/\tau_{jk}^{PRF} \right)^{\sigma} - 1} \right]^{\frac{-\theta+\sigma-1}{\sigma-1}}$, total exports, given by the sum of MFN and preferential exports, can be expressed as:

$$X_{ijk} = \left(\frac{\sigma}{\sigma - 1} \frac{w_i \delta_{ijk}}{P_{jk}} \right)^{1-\sigma} \beta_{jk} E_j M_i \frac{\theta}{\theta - \sigma + 1} \left(\tau_{ijk}^{MFN} \right)^{1-\sigma} \left(\varphi_{ijk}^{MFN} \right)^{-\theta+\sigma-1} \left[1 - s_{ijk} + \left(\frac{\tau_{ijk}^{MFN}}{\tau_{ijk}^{PRF}} \right)^{\sigma-1} s_{ijk} \right] \tag{38}$$

Now substitute the definition of the price index P_{jk} and the MFN cut-off, and re-arrange the last term of (38) to get the gravity equation:

$$X_{ijk} = \frac{\left(w_i \delta_{ijk} \right)^{-\theta} M_i \left(\tau_{ijk}^{MFN} \right)^{1-\frac{\sigma\theta}{\sigma-1}} \left(f_{ijk} w_i \right)^{\frac{-\theta+\sigma-1}{\sigma-1}} \left[1 + \left(T_{ijk}^{\sigma-1} - 1 \right) s_{ijk} \right]}{\sum_{i'} \left(w_{i'} \delta_{i'jk} \right)^{-\theta} M_{i'} \left(\tau_{i'jk}^{MFN} \right)^{1-\frac{\sigma\theta}{\sigma-1}} \left(f_{i'jk} w_{i'} \right)^{\frac{-\theta+\sigma-1}{\sigma-1}} \left[1 + \left(T_{i'jk}^{\sigma-1} - 1 \right) s_{i'jk} \right]} \beta_{jk} E_j \tag{39}$$

Expenditure is defined as in (30), but now tariff revenues account for the share of imports coming under each regime:

$$tr_j = \sum_i \sum_k \left(\frac{\tau_{ijk}^{MFN} - 1}{\tau_{ijk}^{MFN}} (1 - PUR_{ijk}) + \frac{\tau_{ijk}^{PRF} - 1}{\tau_{ijk}^{PRF}} PUR_{ijk} \right) X_{ijk} \tag{40}$$

with $PUR_{ijk} = X_{ijk}^{PRF}/X_{ijk}$.

Finally, in both the eligibility and the effective-tariff models, a country's income is given by its sales revenues:

$$Y_i = \sum_j \sum_k X_{ijk} \tag{41}$$

E.1 Counterfactuals

Consider now the counterfactual exercise of two countries signing an FTA. The tariff applied to the partner country passes from the MFN rate to a preferential rate τ_{ijk}^{PRF}. For the eligible-tariff model, I maintain the standard notation of indicating the base tariffs with τ_{ijk} and the new tariffs with τ'_{ijk}.

Using the exact hat algebra, where $\hat{x} = x'/x$, the new income $Y'_i = \hat{Y}_i Y_i$ is equal to the sum of the new sales $Y'_i = \sum_k \sum_j X'_{ijk}$. For the eligible-tariff model, this gives:

$$\hat{Y}_i Y_i = \sum_k \sum_j \lambda_{ijk} \frac{\hat{\tau}_{ijk}^{1-\frac{\sigma\theta}{\sigma-1}} \left(\hat{w}_i \hat{\delta}_{ijk}\right)^{-\theta} (\hat{w}_i)^{\frac{-\theta+\sigma-1}{\sigma-1}}}{\sum_{i'=1}^{n} \lambda_{i'jk} \hat{\tau}_{i'jk}^{1-\frac{\sigma\theta}{\sigma-1}} \left(\hat{w}_{i'} \hat{\delta}_{i'jk}\right)^{-\theta} (\hat{w}_{i'})^{\frac{-\theta+\sigma-1}{\sigma-1}}} \beta_{jk} \frac{Y_j \hat{Y}_j + D'_j}{1 - \pi'_j} \tag{42}$$

where λ_{ijk} indicates the expenditure share X_{ijk}/E_{jk}. In the standard eligible-tariff model, the new tariff revenue share can be computed as:

$$\pi'_j = \sum_k \beta_{jk} \sum_i \frac{\tau'_{ij,k} - 1}{\tau'_{ij,k}} \hat{\lambda}_{ij,k} \lambda_{ij,k} \tag{43}$$

while the in the effective-tariff model they are given by:

$$\pi'_j = \sum_k \beta_{jk} \sum_i \left[PUR'_{ijk} \frac{\tau_{ij,k}^{\prime PRF} - 1}{\tau_{ij,k}^{\prime PRF}} + \left(1 - PUR'_{ijk}\right) \frac{\tau_{ij,k}^{\prime MFN} - 1}{\tau_{ij,k}^{\prime MFN}} \right] \hat{\lambda}_{ij,k} \lambda_{ij,k}$$

For the effective-tariff model, the change in income is given by:

$$\hat{Y}_i Y_i = \sum_k \sum_j \lambda_{ijk} \frac{\left(\hat{w}_i \hat{\delta}_{ijk}\right)^{-\theta} \left(\hat{\tau}_{ijk}^{MFN}\right)^{1-\frac{\sigma\theta}{\sigma-1}} (\hat{w}_i)^{\frac{-\theta+\sigma-1}{\sigma-1}} \left[1 + \left(T_{ijk}^{\sigma\left(\iota_{ijk}\frac{\theta}{\sigma-1}\widehat{+(1-\iota_{ijk})}\right)-1} - 1\right) s_{ijk}\right]}{\sum_{i'} \lambda_{i'jk} \left(\hat{w}_{i'} \hat{\delta}_{i'jk}\right)^{-\theta} \left(\hat{\tau}_{i'jk}^{MFN}\right)^{1-\frac{\sigma\theta}{\sigma-1}} (\hat{w}_{i'})^{\frac{-\theta+\sigma-1}{\sigma-1}} \left[1 + \left(T_{i'jk}^{\sigma\left(\iota_{i'jk}\frac{\theta}{\sigma-1}\widehat{+(1-\iota_{i'jk})}\right)-1} - 1\right) s_{i'jk}\right]} \beta_{jk} \frac{Y_j \hat{Y}_j + D'_j}{1 - \pi'_j} \tag{44}$$

where

$$1 + \left(T_{ijk}^{\sigma\left(\iota_{ijk}\frac{\theta}{\sigma-1}\widehat{+(1-\iota_{ijk})}\right)-1} - 1\right) s_{ijk} = \frac{1 + \left(\left(T'_{ijk}\right)^{\sigma\left(\iota_{ijk}\frac{\theta}{\sigma-1}+(1-\iota_{ijk})\right)-1} - 1\right) s'_{ijk}}{1 + \left(T_{ijk}^{\sigma\left(\iota_{ijk}\frac{\theta}{\sigma-1}+(1-\iota_{ijk})\right)-1} - 1\right) s_{ijk}}$$

is the change in tariffs due to the FTA, and accounting for the imperfect use of tariff preferences. Note that if a country passes from MFN to an FTA, $T_{ijk} = 1$ as there are no preference in the base state. The parameter ι_{ijk} is a dummy variable that takes the value of one if the PUR is 100%. In this case, $s_{ijk} = 1$ and the model reverts to the eligible-tariff one, where a change in the tariff applies to all trading firms and the elasticity of tariff-inclusive trade is $\sigma\theta/(\sigma - 1) - 1$.

For both models, equations (42) and (44) represent a system of N equations in N unknowns (the national income). By Walras' Law one equation is redundant, hence we fix a numeraire country with $\hat{Y}_i = 1$. Alternatively, we can fix world income as the numeraire.