The dynamics of the trade elasticity: Evidence from the 2018 US trade war

Nicolò Tamberi

July 2025

Centre for Inclusive Trade Policy
Working Paper No.027

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]

Abstract

This paper investigates the trade elasticity, a key parameter in international economics, and examines biases in its estimation arising from dynamic treatment effects and staggered adoption. Leveraging the 2018 US trade war tariffs as a natural experiment, I apply a local projections difference-in-differences approach, estimating a short-run elasticity of -1.4 and a long-run elasticity of -3.7, with adjustments stabilizing within 15 months. Failing to account for staggered adoption or dynamics introduces a downward bias of approximately 50%. I also propose a novel correction to the estimation of dynamic multipliers to cumulative policy changes in the presence of staggered treatment timing. This correction eliminates a systematic bias that can otherwise distort estimates.

Suggested citation

N, Tamberi (2025) The dynamics of the trade elasticity: Evidence from the 2018 US trade war, Centre for Inclusive Trade Policy, Working Paper 027

Non-Technical Summary

How much does trade respond to tariffs? This is a central question for international economics and trade policy, especially at a time when protectionism is on the rise. Economists use the concept of trade elasticity to answer this question, which measures how strongly imports fall when tariff costs go up. This elasticity determines the expected impact of tariff changes on trade flows, and it also plays a key role in calculating the welfare gains from trade in economic models.

While the trade elasticity has been studied extensively, estimates vary widely. Older studies often found large values (around -5), suggesting that trade is highly responsive to tariffs. More recent research, especially using difference-in-differences methods, has reported much lower numbers (around -1.5 to -2.5). But these lower estimates may suffer from biases, which this paper sets out to address.

Two key issues can distort estimates of trade elasticity. First, trade adjusts over time. When new tariffs are introduced, firms often need time to react (renegotiating contracts, switching suppliers, or reorganizing production). This means that short-run responses to tariffs may be smaller than the long-run effects. Second, tariff changes are often implemented in stages. Rather than all products being affected at once, tariffs are typically introduced in waves across different goods and countries. This staggered rollout creates a problem for standard methods. When combined with dynamic responses – that is, gradual adjustments in trade over time – this can lead to a serious underestimation of how strongly trade actually reacts to tariffs.

To study this problem, I focus on the 2018 US trade war, one of the largest and most abrupt shifts in US trade policy in recent history. During this period, the United States imposed new tariffs on many products and countries, often with little warning and using rarely used trade laws. The policy shocks were large, sudden, and varied by product and country, making them well suited for empirical analysis.

Using monthly US imports data by product and country, I apply a modern econometric approach that carefully accounts for both dynamics – how trade responses unfold gradually over time – and staggered adoption – the fact that tariffs were introduced at different times for different products and countries.

The results speak clearly. When using standard methods, the long-run trade elasticity appears to be around -1.9, in line with previous estimates. But once we correct for dynamic responses and the staggered rollout of tariff changes, the long-run elasticity increases in absolute value to -3.7. In other words, standard approaches underestimate the true response by about 50%.

The short-run elasticity is estimated at -1.4, confirming that trade takes time to adjust. However, the adjustment is relatively quick: the elasticity stabilizes after about 15 months. This is notably faster than suggested by previous studies, which often find adjustment periods of five years or more. One possible explanation is that firms perceived the trade war tariffs as temporary, encouraging quicker substitution away from targeted suppliers.

These results have important implications for economic modelling. The magnitude of the trade elasticity feeds directly into calculations of the gains from trade – or how much better off a country is compared to the hypothetical situation of complete self-sufficiency (autarky). Using the lower, biased estimate of -1.9 would imply that the US gains 4.5% in welfare from trade compared to autarky. Using the corrected estimate of -3.7 cuts those gains nearly in half, to 2.4%.

To validate these results, I show that the corrected elasticity estimates accurately predict observed declines in aggregate US imports from China during the trade war, as well as the effects of retaliatory tariffs imposed by the European Union on US products. In both cases, the estimated trade response closely matches what we see in the data.

The data also reveals nearly complete pass-through of tariffs into US import prices. In other words, the entire cost of the new tariffs was borne by US importers, not foreign exporters. However, there is some variation across sectors, with a few showing incomplete pass-through.

The findings have two key takeaways: first, policymakers and analysts using standard trade elasticity estimates may be underestimating the true responsiveness of trade to tariffs. Second, because the empirical setting closely resembles recent US tariff increases, these results can inform how we evaluate current and future protectionist measures.

The Dynamics of the Trade Elasticity: Evidence from the 2018 US Trade War

Nicolò Tamberi*

15th July 2025

Abstract

This paper investigates the trade elasticity, a key parameter in international economics, and examines biases in its estimation arising from dynamic treatment effects and staggered adoption. Leveraging the 2018 US trade war tariffs as a natural experiment, I apply a local projections difference-in-differences approach, estimating a short-run elasticity of -1.4 and a long-run elasticity of -3.7, with adjustments stabilizing within 15 months. Failing to account for staggered adoption or dynamics introduces a downward bias of approximately 50%. I also propose a novel correction to the estimation of dynamic multipliers to cumulative policy changes in the presence of staggered treatment timing. This correction eliminates a systematic bias that can otherwise distort estimates.

1 Introduction

The trade elasticity – the elasticity of imports to variable trade costs – is a foundational parameter in international economics. It governs how countries respond to changes in tariffs and transport costs and plays a central role in both static and dynamic quantitative trade models used to analyse the effects of trade policy. In particular, the trade elasticity shapes the predicted welfare gains from trade liberalization and determines the cross-country transmission of shocks in open-economy macro models. Despite its centrality, there remains considerable uncertainty over its magnitude.

A wide range of estimates exist in the literature. Earlier studies exploiting cross-sectional variation suggest a central value around -5 (Head and Mayer, 2014). More recent studies using

*University of Sussex. I thank L. Alan Winters, Maurizio Zanardi, Chad Bown, participants to the 2025 UKTPO conference in London and the 2024 CITP conference in Glasgow. This research was supported by the ESRC Centre for Inclusive Trade Policy, grant number ES/W002434/1.

difference-in-differences (DiD) frameworks find substantially lower elasticities in the range of -1.5 to -2.5 (Boehm et al., 2023; Amiti et al., 2019; Fajgelbaum, Goldberg et al., 2020). These (absolutely) lower estimates imply that the welfare gains from trade, as calculated using the formula of Arkolakis et al. (2012), are much larger than what would be predicted using the earlier consensus elasticity of -5.

However, difference-in-differences estimates are potentially biased due to two key econometric issues: dynamic treatment effects and staggered adoption. When treatment effects evolve over time, the short- and long-run elasticities may differ substantially. The few studies that account for these dynamics consistently find that the short-run elasticity is significantly lower than long-run elasticity (Boehm et al., 2023; Anderson and Yotov, 2020; Anderson and Yotov, 2025). Second, most empirical applications are characterised by staggered adoption. Staggered adoption occurs when a treatment – e.g., a tariff change – is introduced to different units at different points in time. When combined with dynamic treatment effects, staggered adoption can bias conventional difference-in-differences estimation, necessitating appropriate methodological adjustments.

Staggered adoption introduces a bias because previously treated units are used as control for newly treated ones. When treatment effects are dynamic, the previously treated units can still be experiencing delayed treatment effects. When they are used as control for newly treated ones, we have a bias in the estimation of the treatment effect.

In this paper, I address both concerns and demonstrate their quantitative importance for the estimation of trade elasticity. My findings show that applying difference-in-differences while ignoring dynamics or staggered adoption produces a downward bias of approximately 50% in the absolute value of both short- and long-run trade elasticity. Consequently, the welfare gains from trade predicted by static trade models – evaluated at the long-run elasticity – are roughly half the size of those predicted using an elasticity estimated without accounting for staggered adoption.

My analysis exploits the arguably exogenous variation in tariffs introduced by the 2018 US-initiated trade war as a natural experiment. The 2018 US special tariffs provide valuable econometric variation. First, they vary across countries and products, allowing us to control for unobservable factors with a rich set of fixed effects. Second, they minimise typical endogeneity concerns associated with tariffs. These tariffs were enacted swiftly, leaving little room for anticipatory behaviours, and leveraging rarely-used trade policy tools available to the US president (Bown, 2018a; Bown, 2017). According to Fajgelbaum and Khandelwal (2022), these tariffs represented 'the largest and most abrupt change in US trade policy history'. Bown (2018b) further characterised this shift in trade policy as 'truly different' from the past. The unprecedented and largely unexpected nature of the tariffs ensures confidence in their exogeneity, making them well-suited for estimating tariff elasticity.

Finally, by focussing on a specific instance of tariff variation, I can precisely identify the timing of the tariff changes and examine the surrounding economic environment.

Using monthly US import data at the 10-digit product level, I estimate the elasticity of imports to tariffs using the local projections difference-in-differences (LP-DID) approach introduced by Dube et al. (2023), adapted to a triple-differences setting relating the change in US imports over time across products and exporting countries to changes in tariffs. The estimation covers up to 48 months after a tariff change.

My results show that the short-run elasticity is -1.4 while the long-run elasticity is -3.7, with convergence occurring after roughly 15 months.[1] When staggered adoption is neglected, the short-run elasticity is estimated at -0.66 and the long-run at -1.97 – a 50% downward bias in the absolute value. By contrast, a standard difference-in-differences approach, that does not distinguish between short- and long-run effects or account for the staggered nature of tariff changes, yields an elasticity estimate of -1.86.

These results imply that the long-run gains from trade predicted by static trade models are smaller than we previously thought. In the paper, I show that using an elasticity of -2 implies welfare gains from trade of 4.5% for the US. By using the unbiased estimate of -3.7 instead, the predicted welfare gains are reduced to 2.4%.

A second contribution of this paper is to formally characterise how staggered adoption can affect the estimation of cumulative multipliers (Jordà and Taylor, 2025) using local projections (LP) a technique increasingly used to measure cumulative dynamic treatment effects. I provide a simple and intuitive correction based on a sample selection rule consistent with the LP-DID framework.

In particular, not dealing with staggered adoption may falsely suggest that a policy is mean-reverting, even when the underlying policy changes are one-off and persistent. The corrected LP-IV estimator introduced in this paper eliminates this bias, ensuring that estimated dynamics reflect the actual structure of the treatment. I find that the elasticity of imports to cumulative tariff changes is -3.1 after 36 months from an initial tariff shock. Without applying my correction, the estimated multiplier is biased toward zero by approximately 25%, leading to an understatement of the long-run trade response.

I validate the robustness of the estimated elasticities with two out-of-sample exercises. The first examines aggregate US import declines from China relative to other countries following the trade war. The second exercise focusses on the retaliatory tariffs imposed by the European Union (EU) on US products. In both cases, my elasticity estimates accurately predict the timing and magnitude

[1] This relatively rapid adjustment may be a characteristic of trade wars, a phenomenon discussed in greater detail later in the paper.

of observed trade responses.

I also show that import prices adjust rapidly and exhibit near-complete tariff pass-through, while import values display more persistent dynamics.

Taken together, these results suggest that failing to address staggered adoption and dynamics in trade elasticity estimation can significantly distort both empirical conclusions and policy prescriptions. This paper contributes new evidence, methods, and implications that are relevant across a wide range of applications in international economics and beyond.

The remainder of this paper is organised as follows. Section 2 reviews the literature on trade dynamics and the estimation of the trade elasticity, and Section 3 introduces the methodology used in this paper. Section 4 presents the policy background of the 2018 US special tariffs and Section 5 describes the data used. Section 6 presents the results. A discussion of the results is presented in Section 7 and Section 8 concludes.

2 Literature

To estimate the trade elasticity, researchers use a range of empirical strategies. Some rely on the gravity equation (Head, Mayer and Ries, 2010; Caliendo and Parro, 2015). Others have used a demand system estimation (Feenstra, 1994; Broda and Weinstein, 2006; Ossa, 2015; Imbs and Mejean, 2015; Imbs and Mejean, 2017) or employed simulated method of moments (Simonovska and Waugh, 2014). While most of the available estimates rely on a Constant Elasticity of Substitution (CES) demand framework, some authors explored less restrictive alternatives (Novy, 2013; Chen and Novy, 2022).

Among these approaches, a widely adopted strategy is to exploit variation in import tariffs (Romalis, 2007; Caliendo and Parro, 2015; Amiti et al., 2019; Fajgelbaum, Goldberg et al., 2020; Fontagné et al., 2022). Tariffs represent a variable trade cost that can be measured precisely and offer rich variation across countries, products and time. Other authors used variation in import prices (e.g., Simonovska and Waugh, 2014; Eaton and Kortum, 2002), or in freight costs (Hummels and Schaur, 2013; Hummels, 1999). In this paper I use import tariffs and define trade elasticity specifically as the elasticity of imports with respect to tariff rates.

Much of the earlier empirical literature relied on cross-sectional variation in tariffs across countries and products, abstracting from time dynamics. These papers tend to produce large values of the elasticity in absolute value. In their meta-analysis, Head and Mayer (2014) report a median trade elasticity of –5, based mostly on static cross-sectional designs.

More recent studies apply difference-in-differences techniques exploiting within-country-product tariff changes over time (Amiti et al., 2019; Fajgelbaum, Goldberg et al., 2020). By controlling for time-invariant unobservables, these approaches address a key omitted variable concern emphasized by Baier and Bergstrand (2007), and often produce smaller elasticity estimates around -2. A difference-in-differences approach can strengthen causal claims about the identification of the trade elasticity. However, DiD is potentially subject to biases stemming from dynamics and heterogeneous treatment adoption.

Despite the extensive theoretical research on trade and dynamic adjustments (e.g., Baldwin and Krugman 1989; Dixit 1989; Alessandria and Choi 2021; Alessandria, Choi and Ruhl 2021; Ravikumar et al. 2019), the applied literature has largely overlooked dynamics when estimating trade elasticity. Most trade models based on the gravity equation – the workhorse model for the applied trade economist – are static in nature, or include at most transition dynamics (e.g., Melitz, 2003). This is, perhaps, why applied researcher often neglected dynamics.

Only a handful of recent papers attempt to estimate dynamic trade elasticities. Boehm et al. (2023) use local projections (Jordà, 2005) to estimate the elasticity of trade to a Most Favoured Nation (MFN) tariff changes by comparing MFN to preferential trade (i.e., subject to a Free Trade Agreement, or FTA), and address the potential endogeneity with an IV design focussed on the response of small exporters to MFN tariff changes. Their baseline results point to a one-year horizon elasticity of -0.76 and a long-run elasticity (7-10 years) of -2. Anderson and Yotov (2020) apply a gravity model including a lagged dependent variable on aggregate bilateral trade data (i.e., not by product). Their estimate of the short-run tariff elasticity is -0.54 while the long-run is -4.94. Anderson and Yotov (2025) estimate a dynamic trade elasticity by using country-pair fixed effects covering different intervals of time and find a short-run tariff elasticity of -0.35 and a long-run elasticity of -4.83. The authors consider the regression with the (standard) country-pair fixed effects covering the entire period as the specification yielding the long-run elasticity. Yet this is exactly the specification that came under scrutiny in the difference-in-differences literature (Goodman-Bacon, 2021).[2]

The emerging consensus is that the trade elasticity varies significantly across products (Fontagné et al., 2022), over time (Boehm et al., 2023; Anderson and Yotov, 2020), and possibly across countries (Chen and Novy, 2022). Moreover, the typical dataset used to estimate the trade elasticity is characterised by staggered adoption.

Hence, when estimating the tariff elasticity, the applied economist is confronted with staggered

[2]The difference-in-differences literature focuses on cases where data are indexed by units and time. In the gravity framework with country-level data, the cross-sectional dimension is identified by the exporter-importer pair.

adoption, dynamic and most likely heterogeneous treatment effects. The combination of these three features creates a well-known bias in canonical difference-in-differences fixed-effects estimators (Goodman-Bacon, 2021; De Chaisemartin and d'Haultfoeuille, 2020; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021). The main problem arises because earlier-treated units may serve as controls for later-treated ones. If the treatment effect takes time to fully manifest – as is likely with trade shocks – this structure leads to systematic bias.

I address the issue using the local projections difference-in-differences approach of Dube et al. (2023). The LP-DID method is quite flexible and can easily be applied to a triple-difference setting. It allows for the estimation of dynamic effects with a continuous treatment, as it is the case with tariff changes, while dealing with staggered adoption. Finally, as with the standard local projections method, by estimating separate regressions for each post-treatment horizon, the LP-DID is less computationally expensive than alternative methods dealing with staggered adoption, making it feasible on large datasets as the one used in this paper.

3 Methodology

This section connects the local projections difference-in-differences method to the empirical application to show how the estimation of the trade elasticity can be affected by staggered adoption. Finally, it introduces the concept of staggered adoption in the estimation of elasticities to repeated policy interventions, or the cumulative multiplier. I shows how the estimation of cumulative multipliers with local projections, a common method in macroeconomics, is still affected by staggered adoption. I propose a simple solution based on a sample selection criteria in line with the spirit of LP-DID.

3.1 Local projections difference-in-differences

To introduce the concept of staggered adoption and heterogeneous treatment effects, I adapt the notation of Dube et al. (2023) in line with the empirical application. The outcome variable, $\ln M_{vit}$, is the log of US imports of product v exported by country i at time t. The cross-sectional dimension is defined by a product-country combination vi.

The identifying assumptions are the same for a difference-in-differences regression: no anticipation and parallel trends. As we rely on a triple-difference estimation, the parallel trends assumption says that, in the absence of treatment, the difference between imports of affected and non-affected products would have been the same in affected and non-affected countries.

Following the difference-in-differences literature, assume that the untreated outcome $\ln M_{vit}(0)$ is generated by:

$$E\left[\ln M_{vit}(0)\right] = \alpha_{vi} + \alpha_{vt} + \alpha_{it} \tag{1}$$

where the α_{vi}, α_{vt} and α_{it} are fixed effects at the exporter-product, product-time and exporter-time level.

As is standard in the trade elasticity literature, the continuous treatment variable is measured as the log of one plus the *ad valorem* tariff:

$$\ln \tau_{vit} = \ln(1 + \text{tariff}_{vit}) \tag{2}$$

For instance, a tariff of 10% corresponds to $\ln \tau = \ln(1.1)$. I allow the elasticity of trade with respect to tariffs, $\varepsilon_g(h)$, to vary by treatment cohort g, where a cohort consists of all units treated in the same period p_g. The elasticity can also vary by post-treatment horizon h. Following Dube et al. (2023), the expected value of the outcome variable at time $t+h$ can be expressed as:

$$\begin{aligned} E[\ln M_{vi,t+h}] = \alpha_{vi} + \alpha_{vt+h} + \alpha_{t+h} + \sum_{g=1}^{G} \left[\varepsilon_g(h) \Delta \ln \tau_{vi,t} \times 1(t = p_g)\right] + \\ \sum_{g=1}^{G} \sum_{j=1}^{\infty} \left[\varepsilon_g(h+j) \times \Delta \ln \tau_{vi,t-j} \times 1(t = p_g + j)\right] + \\ \sum_{g=1}^{G} \sum_{j=1}^{h} \left[\varepsilon_g(h-j) \times \Delta \ln \tau_{vi,t+j} \times 1(t = p_g - j)\right] \end{aligned} \tag{3}$$

The first line of equation (3) represents newly treated units at time $p_g = t$. The second line represents the previously treated units, those treated before period t. Finally, the third line represents units with changes in treatment between period $t+1$ and $t+h$.

Define the h-horizon change in a variable $\Delta^h x_t = x_{t+h} - x_{t-1}$. Subtracting $E\left[\ln M_{vi,t-1}\right]$ from

(3) we obtain:

$$
\begin{aligned}
E\left[\Delta^h \ln M_{vit}\right] = (\alpha_{vt+h} - \alpha_{vt-1}) + (\alpha_{t+h} - \alpha_{t-1}) + \sum_{g=1}^{G} \left[\varepsilon_g(h) \Delta \ln \tau_{vi,t} \times 1(t = p_g)\right] + \\
\sum_{g=1}^{G} \sum_{j=1}^{\infty} \left[(\varepsilon_g(h+j) - \varepsilon_g(j-1)) \times \Delta \ln \tau_{vi,t-j} \times 1(t = p_g + j)\right] + \\
\sum_{g=1}^{G} \sum_{j=1}^{h} \left[\varepsilon_g(h-j) \times \Delta \ln \tau_{vi,t+j} \times 1(t = p_g - j)\right] \quad (4)
\end{aligned}
$$

with the α_{vi} differenced-out. Equation (4) clearly shows the sources of bias arising due to staggered adoption and heterogeneous treatment effects. The standard local projections regression of $\Delta^h \ln M_{vit}$ on $\Delta \ln \tau_{vit}$ and fixed effects would only account for the first line of equation (4), but we would have biases arising from lines two and three.

The bias in line two of equation (4) stems from the presence of previously treated units with $\Delta \ln \tau_{vit} = 0$ but $\Delta \ln \tau_{vit-j} \neq 0$ for some $j > 0$. They become control units for those units with $\Delta \ln \tau_{vit} \neq 0$ but might still be subject to delayed treatment effects themselves. Second, when looking at horizon $t + h$ we want to ensure that the control units are not subject to treatment between period t and $t + h$. The LP-DID sample restriction avoids that these kind of comparisons are made.

To help intuition, consider a unit that received a tariff increase in January 2019. If we are estimating the elasticity at a 12-month horizon, we must ensure that this unit did not receive another tariff change between January 2019 and January 2020. Similarly, any control unit must not have received any tariff change before or during this period. This ensures that we are comparing newly treated units to untreated ones, avoiding contamination from lagged treatment effects.

The LP-DID involves the regression:

$$
\Delta^h \ln M_{vit} = \beta_\tau^h \Delta \ln \tau_{vit} + \alpha_{it}^h + \alpha_{vt}^h + \epsilon_{vit} \quad (5)
$$

subject to:

$$
\begin{cases}
\Delta \ln \tau_{vi,t-k} = 0 \,\text{for}\, k \geq -h & \text{control} \\
\Delta \ln \tau_{vi,t} \neq 0 \,\text{and}\, \Delta \ln \tau_{vi,t+k} = 0 \,\text{for}\, k \in (1,h) \text{and}\, \Delta \ln \tau_{vi,t-k} = 0 \,\text{for}\, k \geq 1 & \text{treatment}
\end{cases} \quad (6)
$$

where $\Delta^h \ln M_{vit} = \ln M_{vi,t+h} - \ln M_{vi,t-1}$ and $\ln M_{vi,t}$ is the log of imports of HS 10-digit product v from country i at time t. τ_{vit} is 1 plus the tariff (e.g., for a 10% tariff $\tau = 1.1$), and α_{it}^h and α_{vt}^h are fixed effects at the country-time and product-time level, respectively, and ϵ_{vit} is the error term.

The parameter k determines the stabilization horizon. When we use only not-yet-treated observations as the control group, we set $k = \infty$. If we are willing to assume that treatment stabilizes after k periods, then we can include more observations in the estimation sample by setting k as a finite integer.

By applying (6), the estimation sample is restricted to units that did not experience any change in tariffs up to $t + h$ (first line), and units with a change in tariff at time t but no change in tariff before t or between $t + 1$ and $t + h$ (second line).

The regression model can also be augmented with pre-treatment lags of the outcome variable. In the main specification, I include lags 1 and 3-24 in intervals of 3 (more on this in Section B).

The set of fixed effects means that the estimation strategy is a triple-difference regression: we compare the difference of treated vs non-treated products within a treated country against the same difference in a non-treated country. As treatment is continuous, rather than comparing averages we are comparing slopes.

By contrast, the static triple-difference with three-way fixed effects is:

$$\ln M_{vit} = b^{dd} \ln \tau_{vit} + \alpha_{vi} + \alpha_{it} + \alpha_{vt} + \upsilon_{vit} \tag{7}$$

By estimating only one coefficient, the estimation of model (7) rules out dynamic treatment effects, and can be biased by the presence of staggered adoption.

In Appendix A, I show how the estimation strategy outlined in (5) can be linked to a Melitz (2003)-style model by looking at transition dynamics. In such models, trade responds sluggishly to tariff changes due to hysteresis in the extensive margin: firms that entered under old tariffs may continue exporting even if new conditions would not allow entry. This generates dynamic adjustment paths which increases in magnitude over time.

Defining treatment: It is important to clarify how treatment is defined in this paper. Broadly speaking, one might consider any non-zero MFN or preferential tariff as treatment. However, to leverage arguably exogenous variation in tariffs, this paper defines treatment as the 2018 special tariffs.

In constructing the dataset, I exclude any products affected by tariff changes prior to the onset of the trade war in January 2018, thereby establishing a clean five-year pre-treatment window starting in 2013. This ensures that any earlier tariff fluctuations had fully stabilized before the treatment period began, reducing the risk of confounding delayed treatment effects. In relation to equation (1), the pre-trade war baseline tariff rates are absorbed into the exporter-product fixed effect α_{iv}. These base tariffs determine the expected value of the *levels* of imports, but not their evolution over time.

3.2 Cumulative multipliers with staggered adoption

As mentioned earlier, the recent paper by Boehm et al. (2023) applies the local projections method to estimate the dynamics of the trade elasticity. As the authors address dynamics and endogeneity of tariff changes but not staggered adoption, their estimates offer a valid benchmark for evaluating the importance of staggered adoption.

While equation (4) showed that the standard LP estimator suffers from bias under staggered adoption, the method used by Boehm et al. (2023) is slightly different as it regresses the cumulative change in imports on the cumulative changes in tariffs from period $t-1$ to $t+h$ rather than on the initial tariff change $\Delta \ln \tau_{vit}$. The method is still subject to staggered adoption issues, but requires a different correction compared to the LP-DID framework.

In many settings – including the 2018 US tariffs – treatment is not a one-off intervention but involves a sequence of changes over time (e.g., initial imposition, escalation or rollback). The sample restriction applied by the LP-DID retains only the first observed tariff change for a unit (unless we assume a stabilization horizon). This approach aligns with the standard difference-in-differences framework where treatment is typically absorbing – once a unit is treated, it remains treated permanently.

By contrast, cumulative treatment frameworks estimate the elasticity of outcomes to the total change in treatment over a horizon, which is often the relevant concept in macroeconomic applications (Jordà and Taylor, 2025). We can define the *cumulative multiplier* $m(h)$ at horizon h of an outcome variable y in response to a policy intervention s as:

$$m(h) = \frac{\mathcal{R}^c_{sy}(h)}{\mathcal{R}^c_{ss}(h)} \tag{8}$$

where $\mathcal{R}^c_{sy}(h)$ is the cumulative response of y to a change in the policy s, while $\mathcal{R}^c_{ss}(h)$ is the cumulative change in s between $t-1$ and $t+h$ (see Jordà and Taylor, 2025 for more details). Instead of estimating these two components separately and taking their ratio, $m(h)$ can be calculated in one step with a local projections instrumental variable (LP-IV) regression. This involves running a regression of $\Delta^h y$ on $\Delta^h s$ (the change in policy over the same horizon) while instrumenting $\Delta^h s$ with $\Delta s = s_t - s_{t-1}$ (Ramey and Zubairy, 2018).

Applied to trade elasticity, the LP-IV regression model is:

$$\Delta^h \ln M_{vit} = \beta_\tau^{h,cum} \Delta^h \ln \tau_{vit} + \alpha_{it}^h + \alpha_{vt}^h + \eta_{vit} \tag{9}$$

instrumenting $\Delta^h \ln \tau_{vit} = \ln \tau_{vit+h} - \ln \tau_{vit-1}$ with $\Delta \ln \tau_{vit} = \ln \tau_{vit} - \ln \tau_{vit-1}$.

However, in a panel data with staggered adoption, this method is still susceptible to bias because newly treated units are compared with those already treated. To see this, consider that, at horizon zero, the change in tariff $\Delta \ln \tau_{vit}$ is instrumented with itself. In this case, LP-IV approach is identical to the standard LP regression, which retrieves a biased coefficient.

To link this issue with the one-off treatment case in the LP-DID framework, we can refer back to equation (3). In LP-DID, we are concerned with units being treated before t and units treated between $t+1$ and $t+h$, as these are assigned to the control group despite potential lingering treatment effects. To address this issue, the LP-DID selects only units that are not-yet-treated at $t+h$ and units treated at t but neither before nor again between $t+1$ and $t+h$.

For the cumulative multiplier, however, we need to retain units for which treatment changes at t and again between $t+1$ and $t+h$, as they contribute to the cumulative response. The problem arises with units that are not treated at t but receive treatment between $t+1$ and $t+h$. For these units, $\Delta^h \ln \tau_{vit} \neq 0$ while $\Delta \ln \tau_{vit} = 0$, meaning the first-stage regression of $\Delta^h \ln \tau_{vit}$ on $\Delta \ln \tau_{vit}$ predicts $\Delta^h \ln \tau_{vit} = 0$, effectively assigning them to the control group despite them experiencing some treatment effects.

Intuitively, including units with $\Delta \ln \tau_{vit} = 0$ and $\Delta^h \ln \tau_{vit} \neq 0$ will bias the first-stage regression of $\Delta^h \ln \tau_{vit}$ on $\Delta \ln \tau_{vit}$ toward zero. Consider a hypothetical setting in which each unit experiences only a single tariff change. In the absence of misclassification, this regression would yield a slope of one and an intercept of zero. However, when the sample includes observations with no tariff change at time t but non-zero cumulative changes over the horizon, we are mechanically inflating the intercept while pulling the slope downward. This distortion can create the false impression that tariffs are mean-reverting, even when the underlying policy is persistent. In such cases, the apparent dynamic pattern is not an economic feature of the data but a statistical artefact of improper treatment classification.

To correct for this issue, we have to apply a sample selection to (9). The 'clean' LP-IV regression to retrieve the cumulative multiplier involves the estimation of (9), instrumenting $\Delta^h \ln \tau_{vit}$ with

$\Delta \ln \tau_{vit}$ and subject to:

$$\begin{cases} \Delta \ln \tau_{vi,t-k} = 0 \, \text{for} \, k \geq -h & \text{control} \\ \Delta \ln \tau_{vi,t} \neq 0 \, \text{and} \, \Delta \ln \tau_{vi,t-k} = 0 \, \text{for} \, k \geq 1 & \text{treatment} \end{cases} \tag{10}$$

The control group consists of units not-yet-treated at $t + h$. The treatment group includes units treated at t but not before, allowing for treatment to re-occur between $t + 1$ and $t + h$. Units that are treated for the first time between $t + 1$ and $t + h$ are excluded from the sample.

In Section C of the Appendix I provide Monte Carlo simulations to make the point numerically. The simulations show that the bias is for the standard LP-IV method can be substantial, while the clean version can retrieve the true parameter.

4 Policy background

Starting in 2018, the US imposed a series special tariffs on many products and partners. The tariff increases are related to i) safeguard measures on solar panels and washing machines against most partners; ii) tariffs on steel and aluminium products on the base of national security, imposed against many partners; and iii) tariffs on various Chinese products. The special duties varied from 10% to 50%, and were imposed in different waves, with some countries exempted. Least developed countries were excluded, while other countries negotiated temporary or permanent exemptions.

Apart from the tariffs on solar panels and washing machines, the special tariffs were imposed shortly after they were announced. For instance, the first list of affected Chinese products was announced in April 2018, and tariffs were introduced in July of the same year. Similarly, the second tranche of Chinese tariffs was announced in July 2018 and imposed in September. This short time between announcement and implementation minimises the possibility of large anticipatory behaviours.

The tariffs on metal products added 10% or 25% starting in March 2018 for most countries, covering about $48 billion of imports.[3] Countries negotiated temporary and permanent exemptions from these tariffs, resulting in groups entering and exiting treatment at different point in times.

The safeguard tariffs on solar panels and washing machines were the result of investigations started in spring 2017, and they were implemented in February 2018. Differently from the metals or China tariffs, the tariffs on washer and solar were initially foreseen to last for three and four years,

[3]See the Peterson Institute timeline https://www.piie.com/blogs/trade-and-investment-policy-watch/2018/trumps-trade-war-timeline-date-guide

respectively, but they were subsequently extended.

The tariffs on Chinese products were implemented in three main waves, and represent the largest shock. The first wave added 25% on $50 billion worth of Chinese imports, mainly on machineries and electrical equipment. The second wave targeted $200 billion of imports from China, starting with a 10% additional duty in September 2018 and raised to 25% in May 2019. The second wave covered more consumer products than the first one. Finally, in September 2019, the US imposed a 15% tariff on another $112 billion, hitting mainly clothing and shoes products. They were reduced to 7.5% in February 2020 as part of a deal between the US and China. Figure 1 shows the evolution of the special tariffs over time.

Figure 1: Tariffs waves

(a) China (b) Solar panels and washing machines (c) Metals

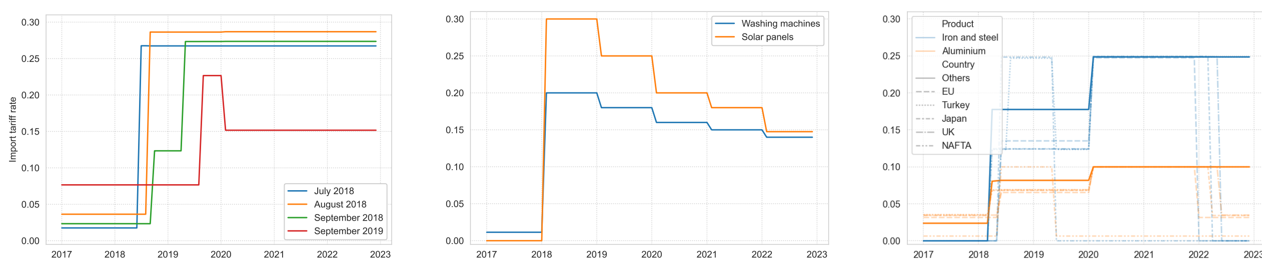

Source: author's elaboration based on USITC data. Panel (a) shows the average tariff on Chinese targeted products, excluding the products which are subject both China tariffs and solar, washer or metal tariffs. Panel (b) plots the average tariff on targeted solar panels and washing machines, excluding exempted countries. Panel (c) shows the average tariff on targeted metal products by country.

Overall, the average tariff applied by the United States increased sharply, passing from 1.77% prior to 2018 to above 4% by the end of 2022.[4]

Figure 2 shows the evolution of import values and prices of affected vs non-affected units (defined as a country-product combination).[5] Panel (a) of Figure 2 shows the evolution of the average import values for affected and non-affected flows, setting the average of 2015 to 100. The two series move closely together until 2018-19, and they start to diverge as tariffs increase. Panel (b) shows the evolution of import prices (in logs) inclusive of the import duties, setting the average of 2015 to

[4] This is computed as trade-weighted average with weights given by US imports over 2013-17 by countries and products.

[5] For illustrative purposes only here, targeted units are identified as units that are treated at some point in time during the sample period. This means that the actual timing of tariff changes is not precise in Figure 2, and also that some of the units treated in 2018 will have seen their tariffs being lifted by the end of the sample period.

zero. As for import values, the two series move closely together up to 2018. As tariff increased in 2018, the import price of affected products increase sharply. While only indicatively, Figure 2 suggests that the effect of tariffs on prices is almost immediate, while it takes more time to fully materialize on import values, a fact that will be confirmed by the econometric analysis.

Figure 2: Targeted and non-targeted products

(a) Import values (b) Tariff-inclusive import prices

Panel (a) shows the average import values of targeted and non-targeted units over time, setting the average of 2015=100. Panel (b) shows the average log import prices (tariff-inclusive) setting the average of 2015 to zero. A unit is defined as a country-product combination, and it is considered affected if it faced additional tariffs at some point over the sample period.

5 Data

Imports data at the 10-digit level of the US product classification by country and with monthly frequency over 2013-22 are taken from the US Census. The tariff-inclusive import prices are computed as the sum of import values and calculated duties over import quantities.

As product classifications were updated during the sample period, I concord products over time using the algorithm developed by Pierce and Schott (2012). I also review the resulting product classification to ensure that products subject to the tariffs are not merged with non-targeted ones.

The tariff data are taken from the USITC website. I only consider *ad valorem* tariffs and drop specific tariffs. In the estimation sample, I only consider tariff changes due to the special tariffs on China, metal products or solar panels and washing machines. This means dropping from the sample country-product series for which either the MFN or the preferential tariff change over time. There are 278 products with changes in the MFN tariffs and 347 country-product pairs with changes in the preferential rates relating to Australia, Japan, and South Korea.

Another issue to take into account is the termination of the Generalised System of Preference (GSP), a trade program granting duty-free access to certain imports from developing countries. As the scheme expired in December 2020, GSP countries saw changes in their market access to the US. I therefore remove from the sample all GSP countries.[6]

Finally, I also remove any unit that is subject to anti-dumping (AD) or countervailing (CV) duties over the sample period. The information on AD and CV duties for the US is taken from the Temporary Trade Measures database (Signoret et al., 2020). I find 9,458 country-product observations subject to AD or CV duties in the sample period. Removing these observations ensures that the control units did not experience any tariff change in the sample period.

For MFN countries, the base pre-treatment tariff is the MFN rate over 2013-15. Because I discard any product subject to changes in the MFN tariff, the 2013-15 rates correspond to the 2013-22 ones. For FTA partners, the base tariff is computed as the log-weighted average of the MFN and preferential tariff, with weights given by the share of imports by tariff program over 2013-15. That is, for a country exporting under the MFN and a preferential tariff, the applied tariff is computed as $\ln \tau_{vit} = s_{vi}^{PRF} \ln \tau_{vit}^{PRF} + (1 - s_{vi}^{PRF}) \ln \tau_{vt}^{MFN}$ where s_{vi}^{PRF} is the 2013-15 share of imports coming under the preferential regime, τ_{vit}^{PRF} is one plus the preferential tariff rate and τ_{vt}^{MFN} is the MFN tariff rate. Such calculation is based on a first order approximation of theoretical model in which firms from an FTA partner country self-select into the preferential or MFN tariff regime (see Tamberi, 2023).

The 2018 special tariffs were additional to the base rate. Hence, when a country-product unit is targeted the extra duty is added on top of the base rate. For an initial rate of 10%, an additional 10% tariff results in a new tariff of 20%. As tariffs were not imposed on the first day of the month, I consider a month treated if the extra tariffs were imposed in the first half of the month. Otherwise, treatment starts the next month.

I follow Fajgelbaum, Goldberg et al. (2020) in not considering special tariffs changes which are subject to quantity threshold. For washing machines, I only apply the 20% tariff and discard the 50% tariff on over-quota units, as well as discarding the 50% on over-quota washing machine parts.

Table 1 reports the summary statistics of the main variables used in the estimation over the period 2013-22. The dataset includes 12.2 million observations with non-zero imports at the country-product level. For tariffs, we do not have missing observations (the full dataset has about 46 million observations), hence we can compute the change in tariffs even if the log of imports is missing. In terms of tariff changes, there are almost 20,000 products with changes related to the tariffs on China and another 19,000 country-product observations seeing changes in terms of metal tariffs. The solar

[6] I also exclude Turkey and India, which were removed from the GSP scheme before its expiration date in 2019.

and washer tariffs contribute less, as they affect only a few products, giving in total 795 changes.

Table 1: Summary statistics

Variable	Count	Mean	Std	Min	Median	Max
Log imports	12,158,907	10.85	2.45	5.53	10.74	22.69
Log (1+tariff)	12,158,907	0.04	0.06	0.00	0.01	1.56
Changes in log(1+tariff)						
China	19,811	0.09	0.07	-0.07	0.12	0.22
Metals	19,395	0.03	0.20	-0.22	0.09	0.22
Solar and washer	795	0.03	0.10	-0.04	-0.02	0.26

The table reports the summary statistics of the two main variables used in the estimation, the log of imports and the log of one plus the applied tariff rate. The summary statics are computed over the period 2013-15. The second panel of the table reports the summary statistics of the changes in the tariff rates by type of tariffs.

6 Results

6.1 Main results

Figure 3 plots the estimated trade elasticity across three methods: the static triple-difference regression of model (7) (dashed line), a standard local projections estimator without adjustment for staggered adoption (orange line), and the LP-DID estimator with proper sample restrictions (blue line). All local projections specifications include lagged log imports (lags 1 and 3-24, in 3-month intervals; see Appendix Section B) and are estimated over horizons from -24 to +48 months relative to the tariff change.

The LP-DID estimates show an immediate elasticity of -1.43, increasing over time in absolute value and reaching -3.74 by horizon 48. The elasticity stabilises approximately by month 15. This quick adjustment time might be a feature of the trade war, a topic discussed further in Section 7.

The comparison with the standard LP estimator (orange line) shows that ignoring staggered adoption leads to a downward bias of roughly 50% in absolute value at all horizons. The static triple-difference estimate (-1.86) closely matches the long-run estimate of the unadjusted LP model (-1.97). If this elasticity is interpreted as a long-run value, we have again an underestimation of about 50%.

Figure 3: Trade Elasticity: Static, LP, and LP-DID Estimates

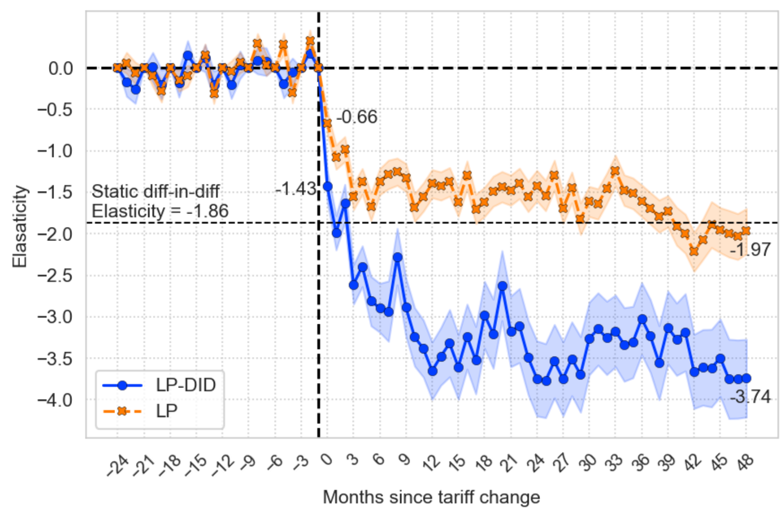

The figure reports the estimated horizon elasticities from the LP and LP-DID estimation together with the 95% confidence interval. Each regression controls for lagged values of the log of imports, including lags 1, 3, 6, 9, 12, 15, 18 and 24. Standard errors are clustered at the country-product level. The dashed horizontal line represents the constant elasticity estimated with the static triple-difference estimator.

These differences are economically significant when calibrating quantitative trade models. Consider the Arkolakis et al. (2012) formula for the gains from trade:

$$G_j = 1 - \lambda_{jj}^{1/\varepsilon} \tag{11}$$

where λ_{jj} is the share of domestic consumption of country j and ε is the absolute value of the (long-run) trade elasticity. A larger value of ε implies smaller gains from trade, as product are more substitutable. As the models considered by ACR are static, the relevant elasticity is the long-run one.

For the United States, Costinot and Rodríguez-Clare (2014) calculates $\lambda_{jj} = 0.913$. With the static estimator ($\varepsilon = 1.87$) we obtain $G_{US} = 4.7\%$. The long-run elasticity of the LP estimator ($\varepsilon = 1.97$) would tell us that the gains from trade for the US are 4.5%. If instead we use the LP-DID value of 3.74, the gains from trade halve at 2.4%.

In short, failing to correct for staggered adoption and dynamics leads to significant overstatement of the welfare gains from trade. Given how central trade elasticity is to structural and policy models, this bias has first-order implications.

6.2 Cumulative multipliers

Figure 4 reports results from the LP-IV estimation of cumulative tariff elasticities. Panel (a) shows the first-stage coefficients (linking the initial tariff change to the cumulative change in tariffs), while Panel (b) shows the second-stage estimates of the cumulative multiplier – i.e., the elasticity of imports to the cumulative tariff change. All regressions control for lagged log imports as in the baseline LP-DID model.

In both panels, I compare two set of estimates. The blue lines represent the 'clean' LP-IV, which applies the sample restriction described in Section 3.2, excluding units treated for the first time between $t+1$ and $t+h$. The orange lines use the 'uncorrected' LP-IV, which does not exclude problematic controls.

Two main findings emerge. First, the second-stage estimates show a significant downward bias from failing to account for staggered adoption. At horizon 36, the uncorrected LP-IV multiplier is 26% lower than the clean estimate. This mirrors the bias seen in the one-off treatment case (Section 6.1), and again results from contaminated comparisons between newly and previously treated units.

Second, we can find marked differences in the first-stage coefficients. As we know the evolution of the 2018 US tariffs, we can link the first-stage results to the pattern of tariff changes. In the clean LP-IV specification, the first-stage coefficient equals 1 for horizons 0-4, reflecting the fact that no unit received a second tariff change within four months of its first. At these horizons, $\Delta^h \ln \tau_{vit} = \Delta \ln \tau_{vit}$ and the clean LP-IV replicates the LP-DID estimates. At longer horizons, we observe shifts in the first-stage coefficient that reflect known features of the 2018-2020 tariff episodes. The dip at horizons 5-6 is due to the third wave of China tariffs, initially imposed in September 2019 and reduced in February 2020. The increase at horizons 7-10 is explained by the second wave of China tariffs, which rose from 10% to 25% after seven months. A drop corresponding to the May 2019 lifting of metal tariffs on Canada and Mexico produces a drop at horizon 11.

On average, an initial tariff is followed by a tariff increase of about 5% from horizon 11 onwards. This result is largely missed using the uncorrected LP-IV method (orange line). Note also that, given that the first-stage coefficients are close to one, the clean cumulative multiplier is close to the elasticity obtained with the LP-DID estimator. This does not need to be the case in other circumstances.

Notably, the estimates from the uncorrected LP-IV estimators for both the first- and second-stage regressions are very close to those obtained by Boehm et al. (2023), who, although using annual data, find a long-run elasticity of -2.12.

Importantly, the pattern observed in the uncorrected first-stage regressions – as well as in the results of Boehm et al. (2023) – may suggest that tariffs are mean-reverting. In the context of the 2018 US tariffs, however, we know this is not the case. As explained in Section (3.2), the apparent mean reversion might not a feature of the data but rather a statistical artefact introduced by misclassification in the estimation sample. The bias arises mechanically when the first-stage regression attempts to predict cumulative tariff changes $\Delta^h \ln \tau$ using only the initial tariff change $\Delta \ln \tau$, while including units for which the latter is zero and the former is non-zero. In such cases, the first-stage regression will underestimate the slope and inflate the intercept. Note that this does not necessarily imply that the tariffs examined in Boehm et al. (2023) are not genuinely mean-reverting, but it is possible that mean reversion is a statistical artefact rather than a feature of the data.

Figure 4: Cumulative multiplier

(a) First-stage results

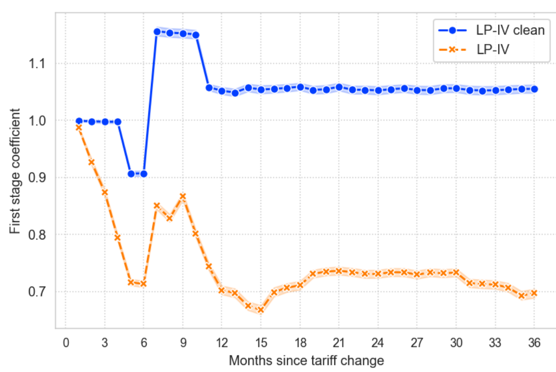

(b) Second-stage results

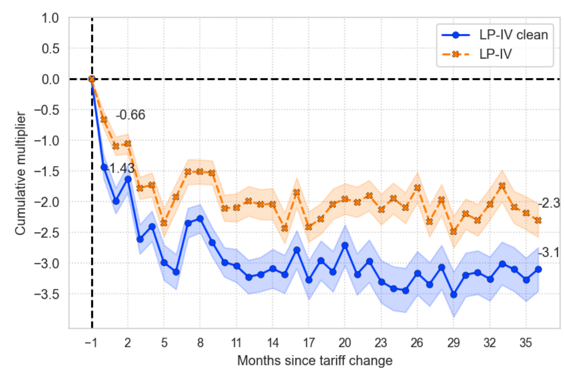

6.3 Validation using external data

To assess the plausibility of the estimated elasticities, I conduct two validation exercises using external data and different levels of aggregation. The first tests whether the LP-DID elasticity presented in Section (6.1) correctly predicts the observed drop in total US imports from China during the trade war. The second examines whether EU imports of products subject to retaliatory tariffs on the US followed dynamics consistent with the estimated elasticity.

A. Total US imports from China: The additional tariffs imposed on China are a shock large enough to be seen on aggregate US goods imports from China. Hence, in this first exercise I compare two estimates for the change in total US imports from China: (i) a model-based prediction using the estimated elasticity and observed product-level tariff changes; and (ii) an empirical benchmark using a difference-in-differences regression comparing US imports from China with other countries' imports from China. While this exercise is not based on truly external data, it provides a sense of how well the elasticities can predict the aggregate changes in imports.

For the prediction, I construct a product-level panel of US imports from China from February 2018 to December 2022. Let $\ln M_v^0$ be the pre-treatment average log US imports of product v from China. I compute the predicted imports h months after the tariff change as:

$$\ln M'_{v,t+h} = \ln M_v^0 + \tilde{\beta}_\tau^h \Delta \ln \tau_{v,t-h} \tag{12}$$

where $\tilde{\beta}_\tau^h$ is the estimated elasticity at horizon h. For repeated tariff changes, effects are added cumulatively. For $h > 48$, the elasticity is held constant at its horizon-48 value, the last estimated horizon.

Predicted imports are obtained $M'_{vt} = \exp(\ln M'_{vt})$, and aggregated across products: $M'_t = \sum_v M'_{v,t}$. The baseline (no-tariff-change) counterfactual is $M^0 = \sum_v \exp(\ln M_v^0)$ and the predicted percentage change is: $100 \times \left(M'_t/M^0 - 1\right)$.

For the empirical benchmark, I use monthly COMTRADE imports from China by 43 countries (OECD + BRICS) over the period January 2017 to December 2022.[7] I estimate the following DiD regression using a Poisson Pseudo Maximum Likelihood (PPML) estimator:

$$M_{jt} = \exp\left[a_{j,month} + a_t + \sum_{t=\text{Feb2018}}^{\text{Dec2022}} b_t \left(US_j \times a_t\right)\right] + \eta_{jt} \tag{13}$$

where $M_{j,t}$ are imports from China of country j at time t, US_j indicates whether the importer is the US. The regression includes importer-by-calendar-month dummies $a_{j,month}$ and time dummies a_t. The DiD coefficient b_t are transformed into percentage changes via $100 \times [\exp(b_t) - 1]$.

B. EU retaliatory tariffs on the US: In June 2018, the EU imposed retaliatory tariffs on a set of US products in response to US metal tariffs. These tariffs – primarily 25% – remained in place until 2022 and were applied only to US exports. In this exercise I check whether LP-DID elasticities correctly predict the evolution of EU imports from the US targeted by the retaliatory tariffs.

The data include monthly imports of the targeted products by EU member states from all major

[7] The countries are: Argentina, Australia, Belgium, Brazil, Bulgaria, Canada, Chile, Croatia, Cyprus, Czechia, Denmark, Estonia, Finland, France, Germany, Greece, Iceland, India, Ireland, Israel, Italy, Japan, Latvia, Lithuania, Luxembourg, Mexico, Netherlands, New Zealand, Norway, Poland, Portugal, Rep. of Korea, Romania, Slovakia, Slovenia, South Africa, Spain, Sweden, Switzerland, Thailand, Turkey, USA and United Kingdom.

exporters (EU, OECD, and BRICS) from January 2017 to December 2021. The estimation is done with the following DiD specification:

$$M_{vijt} = \exp\left[a_{ijv} + a_{jvt} + \sum_{t=\text{Jul2018}}^{\text{Dec2021}} b_t \left(US_i \times a_t\right)\right] \tag{14}$$

where M_{vijt} are imports of product v imported by the EU member j from exporter i at time t. a_{ijv} is an exporter-importer-product fixed effect (the cross-sectional dimension) and a_{jvt} is an importer-product-time fixed effect. US_i is a dummy that takes value of one if the exporter is the US, and it is interacted with time dummies a_t. I retrieve the coefficients b_t and translate them into percentages as $100 \times [\exp(b_t) - 1]$.

I compute the predicted trade path in the same way as in part A, using the LP-DID elasticity and the observed 25% tariff changes. The pre-treatment base value $\ln M_v^0$ is the average log EU import from the US of product v over January 2017 to June 2018. Predicted and baseline values are aggregated as before to compute the percent change in imports.

Figure 5 displays the results for both exercises: aggregate US imports from China (Panel a) and the EU's retaliatory tariffs on US products (Panel b). In each panel, I plot three series: (i) the benchmark PPML difference-in-differences estimates based on external data; (ii) the counterfactual path predicted using the LP-DID elasticities; and (iii) the predictions based on the biased LP elasticities that do not correct for staggered adoption.

In Panel (a), the LP-DID-based predictions track the external DiD estimates remarkably closely, capturing both the timing and magnitude of the import decline following the onset of US tariffs on China. Toward the end of 2022, the external series becomes slightly more negative. This may reflect additional shocks not directly related to tariffs but captured in the empirical series – such as the introduction of the Inflation Reduction Act in August 2022. In contrast, the predictions based on uncorrected LP elasticities consistently underestimate the trade response throughout the period.

Panel (b) tells a similar story for the EU retaliatory tariffs. The LP-DID predictions align closely with the external DiD benchmark across all post-treatment horizons, while the predictions from biased LP elasticities systematically understate the decline in EU imports from the US.

Together, these exercises show that the trade elasticities estimated using the LP-DID approach have strong out-of-sample predictive power, both in terms of dynamic adjustment paths and magnitudes. This supports their external validity, a property rarely assessed in the empirical trade elasticity literature. These results suggest that properly addressing staggered adoption is not only important for identification but also crucial for making accurate policy-relevant predictions.

Figure 5: Validation of results with external data

(a) Predicted change in total US goods imports from China

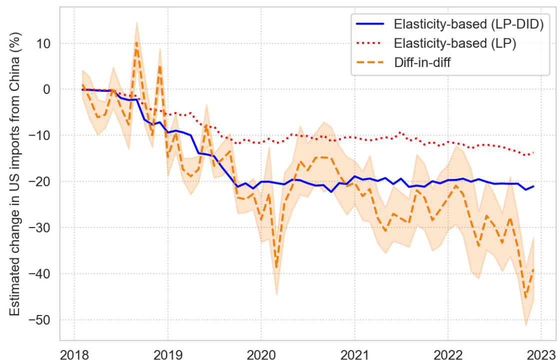

(b) Predicted change in targeted EU imports from the US

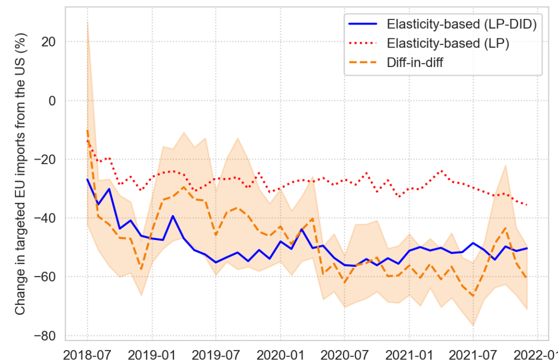

6.4 Results by product groups

Figure 6 reports LP-DID estimates of the trade elasticity by product group, over a 35-month horizon following an initial tariff change. Each regression includes lags 1-12 of the log of imports as control variable.

The adjustment path appears broadly consistent across product categories: in most cases, the elasticity increases in absolute value during the first year and stabilizes around 15 months after treatment – similar to the aggregate baseline result.

Two groups stand out as somewhat more elastic: Chemicals and plastics and Miscellaneous manufactured products, both showing long-run elasticities around -4. In contrast, Machineries have la long-run elasticity closer to -2.5. Metal products are less elastic and exhibit a more volatile response, possibly due to the more complex timing in the metal tariff regime.

For Agrifood, Primary goods, and Textiles and clothing, identification is limited to the second wave of China tariffs, which began with a 10% duty in September 2018 and increased to 25% in May 2019. Since the LP-DID framework focuses on the first treatment, only short-run horizons (0–6 months) are available for these groups.

Overall, while there is some variation in magnitude and volatility, the general shape and timing of adjustment are consistent across sectors. The average of the group-specific elasticities closely tracks the aggregate LP-DID estimate (see Appendix Figure 9).

Figure 6: Results by product groups

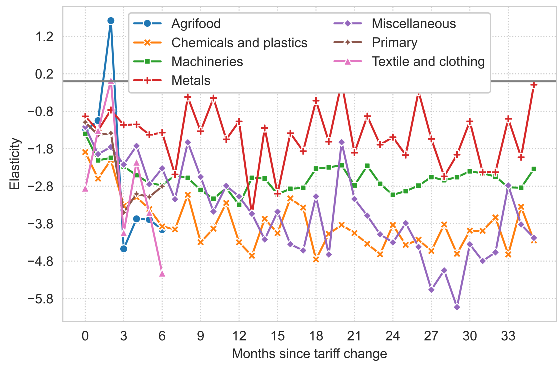

6.5 Asymmetric effects of tariff increases and decreases

I next study whether the response to the imposition of new tariffs or its lifting have asymmetric effects on imports. Using the LP-DID estimator, I run the following regressions:

$$\Delta^h \ln M_{vit} = \beta^h_{\tau+} \Delta \ln \tau_{vit} \times \mathbb{1}(\Delta \ln \tau_{vit} > 0) + \beta^h_{\tau-} \Delta \ln \tau_{vit} \times \mathbb{1}(\Delta \ln \tau_{vit} < 0) + \sum_{l=1}^{12} b^l_M \ln M_{vit-l} + \alpha^h_{it} + \alpha^h_{vt} + e_{vit} \tag{15}$$

where $\mathbb{1}()$ is and indicator function takes the value of one if its argument is true and zero otherwise. Since tariff reductions occurs only after a tariff increase in my dataset, I relax the standard LP-DID sample restriction and assume a stabilization horizon of 10 months. This permits inclusion of reversal episodes without introducing substantial bias, as the dynamic path stabilizes around that point.

Figure 7 Panel a shows the estimated responses to tariff increases and reductions separately. The path following a tariff increase closely matches the baseline LP-DID results, confirming that the relaxed restriction does not distort estimation. The response to tariff removals is somewhat slower in the initial months but converges to the same long-run elasticity, suggesting that the trade response is broadly symmetric over time. The response to a tariff removal appears to be slower at the beginning, but converges to the value of the tariff increase in the long run, with a bump at horizon 23.

Note that the variation for the estimation of the tariff reduction effect beyond horizon 12 come exclusively from the NAFTA exemption from metals tariffs. Other removal events (solar panels, washers, EU/Japan/UK metals tariffs), contribute only to the estimation of horizons 0-12, either because they were phased out in 12-month steps or occurred late in the sample.

Panel (b) shows the hypothetical effect of a tariff increase that is removed after 24 months. In this scenario, we would see trade falling for about 10 months and then stabilising until the additional tariffs are lifted. Then, imports would increase and return back to the pre-tariff increase values in about 13-15 months.

Figure 7: Tariff increases and decreases

(a) Tariff increases and decreases

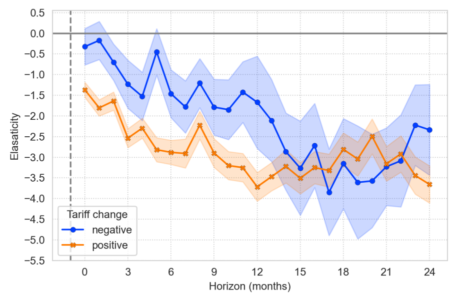

(b) Linked tariff increase and removal

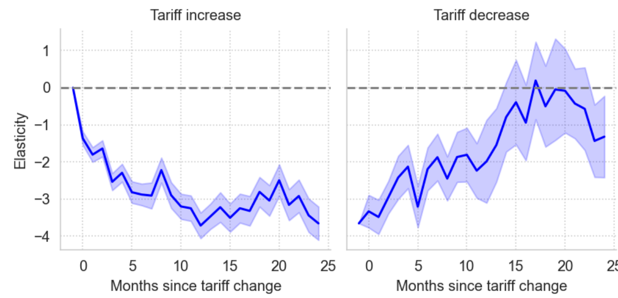

6.6 Prices

Beyond import values and quantities, the US import data also report information on the calculated duties paid at the border.[8] This allows us to construct a measure of tariff inclusive import price as the sum of import value and calculated duty divided by the import quantity $p_{vit} = (M_{vit} + \text{duty}_{vit}) / Q_{vit}$.

I estimate the LP-DID model using log import prices as the outcome, without controlling for pre-treatment trends, since no such dynamics are observed.

Figure 8 show two main results. First, the tariff pass-through to import prices is complete. The average of the post-treatment coefficients is 1.04, confirming the estimates of Amiti et al. (2019) and Fajgelbaum, Goldberg et al. (2020). Second, there adjustment is immediate. The horizons zero and one coefficients are 0.56 and 0.71, respectively, and converge to one by horizon 2. This justifies the use of static estimators in price regressions and explains the consistency of my results with earlier static studies.

Additional results, including a static triple-difference specification by product group, are reported in Appendix Table 2. These show some variation in pass-through across sectors, though the aggregate pattern remains one of near-complete pass-through.

[8]Note that while the calculated duties exclude anti-dumping duties, they do include the 2018 special tariffs.

Figure 8: Tariff-inclusive import prices

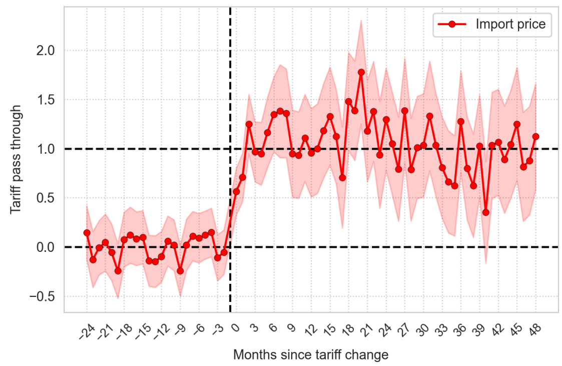

6.7 Additional results and robustness tests

I conduct several robustness exercises to assess the sensitivity of the trade elasticity estimates to sample definition and outcome variable choice. Figure 12 illustrates how progressively removing units affected by other trade policy instruments (e.g., anti-dumping duties, MFN changes, or GSP status) increases the absolute value of the long-run elasticity – from 3.00 using the full sample to the baseline estimate of 3.74 in the cleanest specification.

To address possible distortions from short-lived treatments, I re-estimate the LP-DID model excluding product categories for which only short post-treatment periods are observed (e.g., Agrifood, Primary, and Textiles). I also exclude solar panels and washing machines, whose tariffs may have been anticipated. In both cases (Figure 11a and 11b), the results are very close to the baseline.

Finally, I explore the extensive margin response by estimating a linear probability model in which the dependent variable is an indicator for positive imports. This allows me to capture the role of tariffs in driving products out of the import basket entirely. The results (Figure 10) show a dynamic decline in the probability of importing, with a 1% increase in tariffs reducing the probability of import by 15-20 percentage points over the long run.

7 Discussion

This section discusses the two key empirical results: (i) the short- and long-run trade elasticities, and (ii) the estimated speed of adjustment to tariff shocks. I contextualize both within the broader international economics literature.

Short-run vs long-run elasticity: A long-standing puzzle in international economics is the 'trade elasticity puzzle' – the apparent inconsistency between the trade elasticity values required by International Real Business Cycle (IRBC) models (typically 1-2) and those needed for quantitative trade models (typically 4-8). As emphasized by Ruhl (2008), IRBC models often rely on short-run elasticities, whereas trade models – especially static ones – require long-run values to evaluate welfare gains or counterfactuals.

My results help reconcile this tension. I estimate a short-run elasticity of -1.4, squarely in the IRBC-compatible range, and a long-run elasticity of -3.7, consistent with the lower end of the range used in structural trade models. Although slightly below the Head and Mayer (2014) median of -5, the estimate is robust and grounded in a method that addresses both dynamic treatment effects and staggered adoption, two sources of bias that have affected prior studies.

Moreover, my estimates show that the long-run elasticity is more than twice as large in absolute value as the short-run one – a dynamic pattern consistent with the time-series analysis of Gallaway et al. (2003).

Boehm et al. (2023) report short-run elasticities below 1 (in absolute terms), which is difficult to reconcile with the assumptions of CES-monopolistic competition models. The authors suggest that a new theoretical framework reproducing a short-run elasticity below one might be needed. My estimates, on the other hand, show that a cleaner econometric exercise might be enough to solve this puzzle.

Adjustment speed: A second key result is that imports respond to new tariffs relatively quickly, stabilizing within 15 months. This is notably faster than estimates from other settings: Boehm et al. (2023) find a 7-year adjustment horizon for MFN tariffs, and Egger et al. (2022) estimate 5-10 years for trade agreements to fully take effect.

The shorter adjustment horizon in my setting may reflect key differences in the nature of the policy shocks. The 2018 US special tariffs were introduced unilaterally and often under legal mechanisms that allowed for exemptions and reversals. This legal and political context may have led firms to perceive the measures as temporary, encouraging rapid disinvestment or substitution with the expectation that tariffs might be lifted – a pattern supported by the symmetric adjustment results shown in Section 6.5.

Furthermore, the country-specific targeting of the trade war, especially against China, likely facilitated reallocation. Since the tariffs were not global in scope, firms could more easily switch to alternative suppliers. This is consistent with evidence from Flaaen et al. (2020), who document similar relocation dynamics in the context of U.S. anti-dumping duties. In contrast, MFN tariff changes affect all trading partners simultaneously, reducing the scope for substitution and potentially requiring deeper supply chain restructuring – processes that are slower to unfold. Trade agreements, while typically liberalizing, also differ substantially: they tend to be phased in gradually and increase long-run certainty for firms, which may prompt major reorganization of trade and production networks that takes longer to implement.

Notably, the shorter adjustment period I document is in line with findings from the anti-dumping literature. Sandkamp (2020), using annual data, shows that import values and prices adjust to anti-dumping duties within two years – again reflecting the temporary nature of these policy instruments. Taken together, these observations suggest that the expected duration and scope of trade policy changes play a crucial role in determining how quickly trade flows adjust. Whether this generalizes across settings remains an open question, but the evidence here points to a potentially systematic difference between temporary and permanent trade interventions in shaping adjustment dynamics.

8 Conclusions

The paper showed how ignoring the dynamics and staggered adoption in the estimation of the tariff elasticity can lead to severe downward biases up to 50% on the absolute value of the elasticity. Using a local projections difference-in-differences estimator that addresses both issues, I find that the short-run tariff elasticity is -1.4 while the long-run is at -3.7.

These estimates are externally validated using aggregate US import data from China and EU imports of US goods targeted by retaliatory tariffs. In both cases, predictions based on the estimated elasticities closely match observed trade patterns in terms of both magnitude and timing.

Beyond the trade elasticity estimates, the paper contributes a corrected LP-IV estimator for cumulative multipliers under staggered adoption, which is a tool that can be applied in a wide range of empirical settings when policies are implemented gradually. I show that, as with the one-off treatment elasticity, failing to account for heterogeneity in treatment timing biases the estimated multiplier downward.

In terms of the adjustment path of imports to tariff, I find that most of the action occurs in the first 15 months following a tariff change. This short adjustment period might be a feature of trade wars, where tariff changes might be perceived to be temporary. Moreover, as most of the variation used to identify the elasticity comes from tariffs imposed on China, relocation of production to other countries might explain the quick adjustment periods.

More broadly, these findings suggest that properly accounting for staggered and dynamic treatment effects is essential for recovering credible estimates of trade elasticities and for interpreting their implications in policy and structural model calibration.

References

Alessandria, George and Horag Choi (2021). "The dynamics of the US trade balance and real exchange rate: The J curve and trade costs?" In: *Journal of International Economics* 132, p. 103511.

Alessandria, George, Horag Choi and Kim J Ruhl (2021). "Trade adjustment dynamics and the welfare gains from trade". In: *Journal of International Economics* 131, p. 103458.

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.

Anderson, James E and Yoto V Yotov (2020). "Short run gravity". In: *Journal of International Economics* 126, p. 103341.

— (2025). "Gravity in transition". In: *Review of International Economics*.

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

Baier, Scott L and Jeffrey H Bergstrand (2007). "Do free trade agreements actually increase members' international trade?" In: *Journal of International Economics* 71.1, pp. 72–95.

Baldwin, Richard and Paul Krugman (1989). "Persistent trade effects of large exchange rate shocks". In: *The Quarterly Journal of Economics* 104.4, pp. 635–654.

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.

Bown, Chad P (2017). "Trump's threat of steel tariffs heralds big changes in trade policy". In: *Washington Post* 21.

— (2018a). "Donald Trump's solar and washer tariffs may have now opened the floodgates of protectionism". In: *Peterson Institute for International Economics*.

— (2018b). "Trump has announced massive aluminum and steel tariffs. Here are 5 things you need to know". In: *The Washington Post, March* 1.

Broda, Christian and David E Weinstein (2006). "Globalization and the Gains from Variety". In: *The Quarterly Journal of Economics* 121.2, pp. 541–585.

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 and Pedro HC Sant'Anna (2021). "Difference-in-differences with multiple time periods". In: *Journal of Econometrics* 225.2, pp. 200–230.

Chen, Natalie and Dennis Novy (2022). “Gravity and heterogeneous trade cost elasticities”. In: *The Economic Journal* 132.644, pp. 1349–1377.

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.

De Chaisemartin, Clément and Xavier d’Haultfoeuille (2020). “Two-way fixed effects estimators with heterogeneous treatment effects”. In: *American Economic Review* 110.9, pp. 2964–2996.

Dixit, Avinash (1989). “Entry and exit decisions under uncertainty”. In: *Journal of Political Economy* 97.3, pp. 620–638.

Dube, Arindrajit et al. (2023). *A local projections approach to difference-in-differences event studies*. Tech. rep. National Bureau of Economic Research.

Eaton, Jonathan and Samuel Kortum (2002). “Technology, geography, and trade”. In: *Econometrica* 70.5, pp. 1741–1779.

Egger, Peter H, Mario Larch and Yoto V Yotov (2022). “Gravity estimations with interval data: Revisiting the impact of free trade agreements”. In: *Economica* 89.353, pp. 44–61.

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

Fajgelbaum, Pablo D and Amit K Khandelwal (2022). “The economic impacts of the US–China trade war”. In: *Annual Review of Economics* 14.1, pp. 205–228.

Feenstra, Robert C (1994). “New product varieties and the measurement of international prices”. In: *American Economic Review*, pp. 157–177.

Flaaen, Aaron, Ali Hortaçsu and Felix Tintelnot (2020). “The production relocation and price effects of US trade policy: the case of washing machines”. In: *American Economic Review* 110.7, pp. 2103–2127.

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

Gallaway, Michael P, Christine A McDaniel and Sandra A Rivera (2003). “Short-run and long-run industry-level estimates of US Armington elasticities”. In: *The North American Journal of Economics and Finance* 14.1, pp. 49–68.

Goodman-Bacon, Andrew (2021). “Difference-in-differences with variation in treatment timing”. In: *Journal of Econometrics* 225.2, pp. 254–277.

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.

Hummels, David L (1999). "Toward a geography of trade costs". In: *GTAP Working Paper*.

Hummels, David L and Georg Schaur (2013). "Time as a trade barrier". In: *American Economic Review* 103.7, pp. 2935–2959.

Imbs, Jean and Isabelle Mejean (2015). "Elasticity optimism". In: *American Economic Journal: Macroeconomics* 7.3, pp. 43–83.

— (2017). "Trade elasticities". In: *Review of International Economics* 25.2, pp. 383–402.

Jordà, Òscar (2005). "Estimation and inference of impulse responses by local projections". In: *American Economic Review* 95.1, pp. 161–182.

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

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

Novy, Dennis (2013). "International trade without CES: Estimating translog gravity". In: *Journal of International Economics* 89.2, pp. 271–282.

Ossa, Ralph (2015). "Why trade matters after all". In: *Journal of International Economics* 97.2, pp. 266–277.

Pierce, Justin R and Peter K Schott (2012). "A concordance between ten-digit US harmonized system codes and SIC/NAICS product classes and industries". In: *Journal of Economic and Social Measurement* 37.1-2, pp. 61–96.

Ramey, Valerie A and Sarah Zubairy (2018). "Government spending multipliers in good times and in bad: evidence from US historical data". In: *Journal of Political Economy* 126.2, pp. 850–901.

Ravikumar, B, Ana Maria Santacreu and Michael Sposi (2019). "Capital accumulation and dynamic gains from trade". In: *Journal of International Economics* 119, pp. 93–110.

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

Ruhl, Kim J (2008). "The international elasticity puzzle". In: *unpublished paper, NYU*, pp. 703–726.

Sandkamp, Alexander (2020). "The trade effects of antidumping duties: Evidence from the 2004 EU enlargement". In: *Journal of International Economics* 123, p. 103307.

Signoret, Jose et al. (2020). *Temporary Trade Barriers Database including the Global Antidumping Database.*

Simonovska, Ina and Michael E Waugh (2014). "The elasticity of trade: Estimates and evidence". In: *Journal of International Economics* 92.1, pp. 34–50.

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.
Tamberi, Nicolò (2023). "Preferential trade agreements under uncertainty". In: *Centre for Inclusive Trade Policy Working Paper* 007.

A Connection to theory

In this section I show how the empirical specification of Section 3.1 can be derived by looking at the transition dynamics of a theoretical trade model with heterogeneous firms. In this type of model, the sluggish adjustment of trade values to a change in the tariff rate is due to the extensive margin response. Firms incurs in an upfront sunk cost to start exporting and exit only if hit by an exogenous exit shock. As a consequence, when entry conditions change (e.g., because of a tariff change) firms that were already exporting continue to do so, even if they would have not entered under the new conditions. These legacy firms create hysteresis in the value of trade, which reaches the new equilibrium only after the extensive margin fully adjusted.

As the general settings are quite standard to trade models, I will only introduce them briefly and focus on the transition dynamics. Consumers in country j have CES preferences over products v produced in country i. Firms operate in monopolistic competition, are heterogeneous in productivity φ and produce using only labour. The firm receives a price p_{vijt} from selling q_{vijt} units, hence we have revenues:

$$r_{vijt}(\varphi) = (p_{vijt}(\varphi))^{1-\sigma}\, \tau_{vijt}^{-\sigma} Y_{jt} P_{jt}^{\sigma-1}$$

where the price $p_{vijt}(\varphi) = \frac{\sigma}{\sigma-1}\frac{w_{it}\delta_{vij}}{\varphi}$ takes the form of constant markup (denoted by $\frac{\sigma}{\sigma-1}$) over marginal costs w_{it} and δ_{vij} is an iceberg trade cost.[9] The tariff τ_{vijt} generates a gap between the price paid by the consumer and the one received by the firm. To export, a firm pays an initial sunk cost f_{vij}, hence only firms with a productivity above the cut-off φ^*_{vijt} start exporting. Setting the present value of revenues equal to the fixed cost of exporting we can find the productivity cut-off φ^*_{vijt}. Integrating firm sales $r_{vijt}(\varphi)$ over productivity yields the aggregate imports of product v of country j exported by i.

For total imports observed at time t, we have to take into account both firms that entered in period t facing the current entry cut-off as well as firms which entered in $t-l$ with an entry cut-off φ^*_{vijt-l} and survived, discounting them by the survival probability γ. We can express aggregate

[9] To simplify the derivation, I assumed that variable trade costs δ_{vij} are constant over time. For the empirical application, if variable trade costs are time-varying, identification requires tariffs to be orthogonal to other variable trade costs.

imports at time t:

$$M_{vijt} = \underbrace{\tau_{vijt}^{-\sigma} Y_{jt} P_{jt}^{\sigma-1} \left(\frac{\sigma}{\sigma-1} w_{it} \delta_{vij} \right)^{1-\sigma}}_{\text{Intensive margin}} \underbrace{\sum_{l=0}^{\infty} \gamma^l m_{i,t-l} \left[\int_{\varphi^*_{vijt-l}}^{\infty} \varphi^{\sigma-1} dG(\varphi) \right]}_{\text{Ext. margin active at time } t} \tag{16}$$

The intensive margin is composed by the mass of firms in the exporting country $m_{i,t-l}$ and the integral of their productivity, computed for firms above the exporting cut-off φ^*_{vijt-l}. As firms that entered before period t can still be active if they entered before t and were not hit by the exogenous exit shock $(1-\gamma)$, the extensive margin active at time t takes into account the discounted lagged entry.

The exporting productivity cut-off can be derived by setting expected export revenues equal to the fixed cost of exporting, and solving for productivity. As a simplification, assume that firms expect future values to be the same as current ones. Hence the expected value of exporting is given by:

$$V_{vijt}(\varphi) = \sum_t \gamma^t r_{vijt}(\varphi) = \frac{r_{vijt}(\varphi)}{1-\gamma} \tag{17}$$

Setting $V_{vijt}(\varphi) = f_{vij}$ and solving for productivity yields the exporting cut-off:

$$\varphi^*_{vijt} = \left[\frac{(1-\gamma) f_{vij}}{\tau_{vijt}^{-\sigma} Y_{jt} P_{jt}^{\sigma-1} (w_{it}\delta_{vij})^{1-\sigma}} \right]^{\frac{1}{\sigma-1}} \tag{18}$$

Assuming a Pareto distribution for firm productivity $G(\varphi) = 1 - \varphi^{-\kappa}$ as customary in the literature, and solving the integral in (16) yields the following expression for the value of imports at time t:

$$M_{vijt} = \tau_{vijt}^{-\sigma} Y_{jt} P_{jt}^{\sigma-1} \left(\frac{\sigma}{\sigma-1} w_{it} \delta_{vij} \right)^{1-\sigma} \frac{\kappa}{\kappa-\sigma+1} \sum_{l=0}^{\infty} \gamma^l m_{i,t-l} \left(\varphi^*_{hijt-l} \right)^{-\kappa+\sigma-1} \tag{19}$$

To derive an empirical equation, I plug (18) into (19) and take a first-order log-linear approximation of all time-varying quantities around a steady state level ss:

$$\begin{gathered}\ln M_{vijt|ss} = \ln M_{vij,ss} + \\ \sum_{l=0}^{\infty} \frac{\partial \ln M_{vijt}}{\partial \ln A_{jt-l}} \ln \left(\frac{A_{jt-l}}{A_{j,ss}} \right) + \sum_{l=0}^{\infty} \frac{\partial \ln M_{vijt}}{\partial \ln w_{it-l}} \ln \left(\frac{w_{it-l}}{w_{j,ss}} \right) + \sum_{l=0}^{\infty} \frac{\partial \ln M_{vijt}}{\partial \ln m_{i,t-l}} \ln \left(\frac{m_{i,t-l}}{m_{i,ss}} \right) + \\ \sum_{l=0}^{\infty} \frac{\partial \ln M_{vijt}}{\partial \ln \tau_{vij,t-l}} \ln \left(\frac{\tau_{vij,t-l}}{\tau_{vij,ss}} \right)\end{gathered} \tag{20}$$

where $A_{jt} = Y_{jt} P_{jt}^{\sigma-1}$. Line one of (20) represents the steady state level of log imports, that can be absorbed by fixed effects. Note also that, at any time t, the terms in line two are constant at the jt and it level hence they can be absorbed by fixed effects. Expression (20) can then be written compactly as:

$$\ln M_{vijt} = \alpha_{vij} + \alpha_{jt} + \alpha_{it} + \sum_{l=0}^{\infty} \beta_l \ln \tau_{vijt-l} + \epsilon_{vijt} \tag{21}$$

where the αs summarise aggregate conditions, β_l represent the elasticity of imports to the tariff rate of $t-l$ periods ago (effectively a combination of the structural parameters σ, κ and the discount factor γ), and ϵ_{vijt} is the reminder of the approximation. The short-run elasticity is given by β_0 while the long-run elasticity is given by $\varepsilon^{LR} = \sum_{l=0}^{\infty} \beta_l$. In connection to local projections, the h-horizon elasticity can be computed as $\varepsilon^h = \sum_{l=0}^{h} \beta_l$.

In terms of structural parameters, consider that the partial derivative of the log-imports with respect to past tariffs is given by:

$$\frac{\partial \ln M_{vijt}}{\partial \ln \tau_{vij,t-l}} = \frac{\gamma^l m_{i,t-l} \left(\frac{w_{it-l}^{\sigma-1} \tau_{vijt}^{\sigma}}{A_{jt-l}} \right)^{\frac{-\kappa+\sigma-1}{\sigma-1}}}{\sum_{l=0}^{\infty} \gamma^l m_{i,t-l} \left(\frac{w_{it-l}^{\sigma-1} \tau_{vijt}^{\sigma}}{A_{jt-l}} \right)^{\frac{-\kappa+\sigma-1}{\sigma-1}}}$$

which evaluated at the steady state yields:

$$\frac{\partial \ln M_{vijt|ss}}{\partial \ln \tau_{vij,t-l}} = -\sigma \frac{\kappa - \sigma + 1}{\sigma - 1} \gamma^l (1-\gamma)$$

It is easy to show that the long-run elasticity is given by $\varepsilon^{LR} = -\frac{\sigma\kappa}{\sigma-1}$, which corresponds to the elasticity that we would get from a static model. The interpretation of the current tariff elasticity β_0 in equation (21) depends on whether we assume that the extensive margin responds immediately to a change in tariffs or with a one-period lag (as done by Boehm et al., 2023). In the latter case, the immediate elasticity is simply the intensive margin elasticity and we have $\beta_0 = -\sigma$. If instead we assume that the extensive margin responds immediately to tariffs, we have $\beta_0 = -\sigma - \sigma\frac{\kappa-\sigma+1}{\sigma-1}(1-\gamma)$. These considerations do not affect the interpretation of the long-run elasticity.

To connect theory with the local projections estimation, consider a one-off change in tariffs at time t, with tariffs being constant at $\tau_{vij,0}$ before then. Imports before the tariff change are at their long-run equilibrium:

$$\ln M_{vij,t-1} = \alpha_{vij} + \alpha_{jt-1} + \alpha_{it-1} + \ln\tau_{vij,0}\sum_{l=0}^{\infty}\beta_l + \epsilon_{vijt-1} \tag{22}$$

At time t, tariffs move from $\tau_{vij,0}$ to $\tau_{vij,1}$. Then h periods after the tariff change we have:

$$\ln M_{vij,t+h} = \alpha_{vij} + \alpha_{jt+h} + \alpha_{it+h} + \ln\tau_{vij,1}\sum_{l=0}^{h}\beta_l + \ln\tau_{vij,0}\sum_{l=h+1}^{\infty}\beta_l + \epsilon_{vij,t+h} \tag{23}$$

Taking the difference between (23) and (22) we obtain:

$$\Delta^h \ln M_{vijt} = \alpha_{jt}^h + \alpha_{it+h}^h + \ln\tau_{vij,1}\sum_{l=0}^{h}\beta_l + \ln\tau_{vij,0}\sum_{l=h+1}^{\infty}\beta_l + \epsilon_{vijt}^h - \ln\tau_{vij,0}\sum_{l=0}^{\infty}\beta_l$$

where $x_t^h = x_{t+h} - x_{t-1}$. We can split the last sum into two parts as can be written as $\sum_{l=0}^{\infty}\beta_l = \sum_{l=0}^{h}\beta_l + \sum_{l=h+1}^{\infty}\beta_l$. Rearranging, we obtain:

$$\Delta^h \ln M_{vijt} = \alpha_{jt}^h + \alpha_{it+h}^h + (\ln\tau_{vij,1} - \ln\tau_{vij,0})\sum_{l=0}^{h}\beta_l + \epsilon_{vijt}^h \tag{24}$$

as $\Delta\ln\tau_{vij,t} = \ln\tau_{vij,1} - \ln\tau_{vij,0}$, equation (24) shows the connection of the empirical LP-DID regression model with the transition dynamics predictions of a standard trade model with heterogeneous firms.

B Additional results

Figure 9: Baseline results and average of product groups elasticities

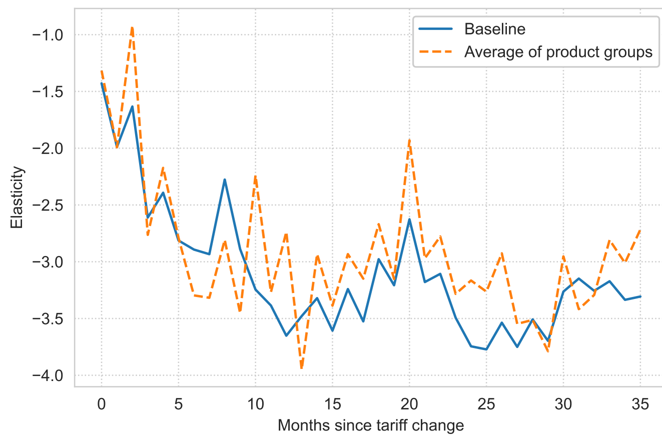

The figure shows the baseline elasticity against the average of the product groups elasticities.

Figure 10: Probability of importing

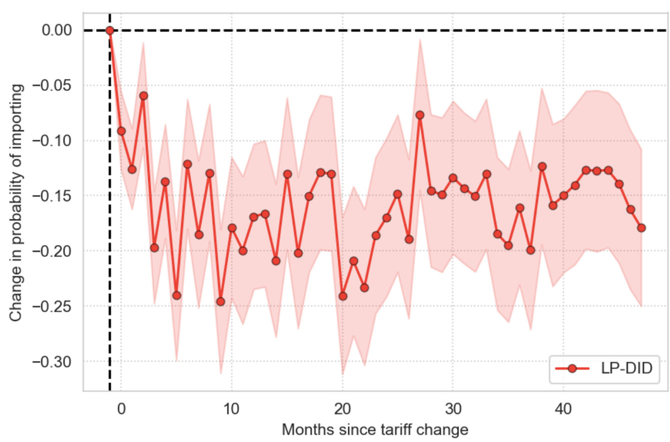

The figure shows the LP-DID estimation of the linear probability model for the probability of importing. The dependent variable is an indicator that equal one if imports are positive and zero otherwise.

Figure 11: Robustness tests

(a) Excluding Agrifood, Primary and Textile and Clothing

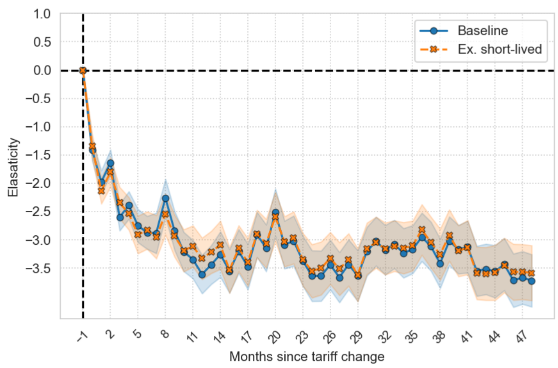

(b) Excluding solar panels and washing machines

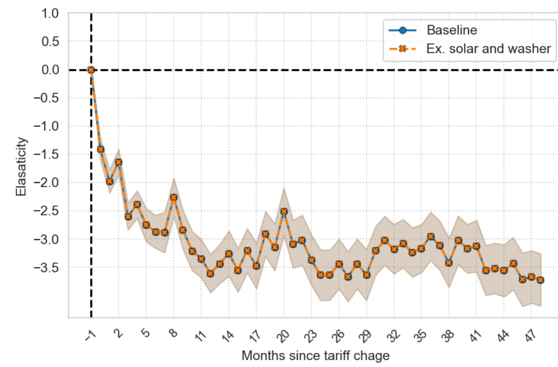

Sub-figure 11a shows the elasticity estimated excluding Agrifood, Primary and Textile and Clothing products together with the baseline results. For those three product groups, we could estimate the elasticity only up to horizon 6 as they then saw subsequent changes in the tariff rates. Sub-figure 11b shows the elasticity estimated excluding solar panel and washing machines together with the baseline results.

Figure 12: Cleaning the sample from other tariff changes

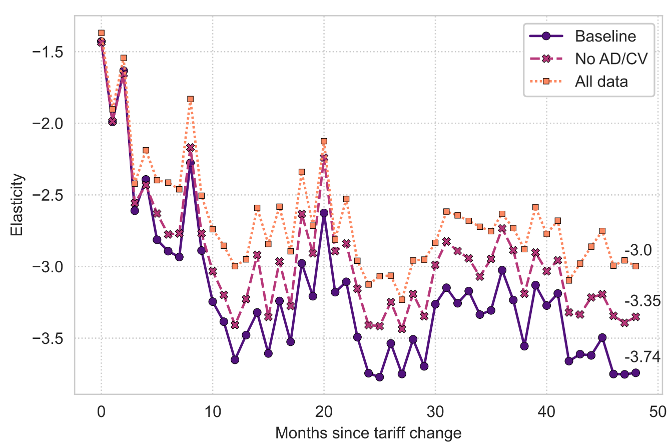

The figure shows the results for the tariff elasticity estimated with the LP-DID on different samples. The line dotted line with squares as markers uses all data. The dashed lines with x as markers removes country-product units subject to AD or CV duties. The solid line with circles as markers is the baseline, and removes also units with changes in the MFN or preferential tariff, as well as GSP countries.

Pre-trends The baseline specification controls pre-treatment values of the outcome variable. The choice is based on an empirical investigation of pre-trends. To check for the existence of pre-trends, I estimate the LP-DID regression from horizon -24 to horizon -2 without controlling for pre-treatment values of imports. The regression for the pre-treatment horizon is:

$$\ln M_{vit+h} - \ln M_{vit-1} = b_\tau^h \Delta \ln \tau_{vit} + a_{it} + a_{vt} + e_{vit}$$

for $h = -24, ..., -2$. Hence the dependent variable for horizon -24 is the change $\ln M_{vit-24} - \ln M_{vit-1}$. As growth rates are easier to interpret if they follow the natural order of time, I report the *negative* of the estimated coefficient corresponding to $\ln M_{vit-1} - \ln M_{vit-24}$, the change from $t-24$ to $t-1$.

I then add as control variables the lags 1-12 of the log of imports, as well as the lags 1, 3, 6, 9, 12, 15, 18 and 24, which are those used for the baseline result. These lagged outcome variables are measured before treatment. More importantly, in the LP regression pre-treatment outcome values are pre-determined, hence the standard issues of lagged dependent variable in panel data do not apply.

Results are plotted in Figure 13. Without controlling for lagged imports, I find some evidence of pre-trends. The horizons -24 to -6 are generally positive and significant. Controlling for the lags 1-12 removes the trend over horizons -12 to -2 mechanically, but we still see some pre-treatment trends from horizons -24 to -13. On the other hand, controlling for lags 1 and 3-24 in intervals of 3 essentially removes all signs of pre-trends. Controlling for lagged values of imports increases the absolute value of the elasticity, in particular from horizon 22 onwards.

Figure 13: Controlling for lagged imports

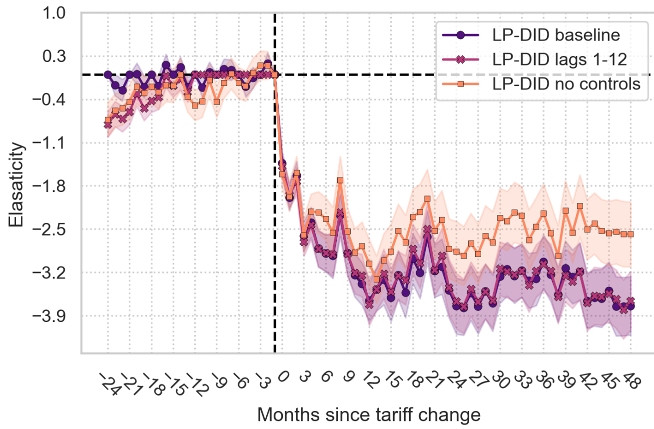

The figure reports the comparison of the elasticity from using the LP-DID estimator with and without the controlling for lagged imports. The LP-DID pre (a) includes lags 1-12 of the log of imports. The LP-DID pre (b) includes lags 1 and 3-24 in intervals of 3 of the log of imports. The 95% confidence intervals are based on standard errors clustered at the exporter-product level.

Price regressions by product: I estimate a static triple-difference regression with exporter-product, exporter-time and importer-time fixed effects. I then interact the log of one plus the tariff with dummies for each product groups. Compared to the LP-DID, the static regression gains in efficiency.

The results are reported in Table 2. On aggregate, the tariff coefficient is 1.014 (0.048), indicating full pass-through (see column 1). There is however some heterogeneity across product groups. First, we note that the average of the product-specific coefficients weighted by the number of observations in each product group yields 1.047. For individual product groups, we cannot reject the hypothesis of full pass-through for Machineries, Miscellaneous manufacturing products, Textile and clothing as well as the residual category. On the other hand, Agrifood, Metals and Primary products exhibit incomplete pass-through. Chemicals and plastics is the only product group for which we cannot reject the null hypothesis of a more than full pass-through.

Table 2: Import price regressions

	(1)	(2)
Aggregate	1.014***	
	(0.048)	
Agrifood		0.680***
		(0.092)
Chemicals and plastics		1.571***
		(0.090)
Machineries		1.154***
		(0.090)
Metals		0.766***
		(0.060)
Miscellaneous		1.084***
		(0.217)
Primary		0.555***
		(0.109)
Textile and clothing		1.015***
		(0.077)
Residual category		2.236*
		(1.339)
Product-time	x	x
Country-time	x	x
Country-product	x	x
Observations	8,661,772	8,661,772
R^2	0.869	0.869

Standard errors clustered at the country-product level in parenthesis. * p<0.1, ** p<0.05, *** p<0.01. The table reports the results from the triple-difference regressions for import prices. The dependent variable is the log of the tariff-inclusive import price. In column 1, the independent variable is the log of one plus the tariff. In column two, the log of one plus the tariff is interacted with dummies for each product group.

C Monte Carlo simulations for clean multipliers

This section presents the Monte Carlo simulations for the estimation of multipliers using a clean control group. The control group consists of units not yet treated at time $t+h$. The treatment groups is composed by units treated at t and possibly re-treated between $t+1$ and $t+h$. Units treated between $t+1$ and $t+h$ but not at t are excluded from the sample.

The data generating process is the following. The panel dataset is composed $N = 100$ units and $T = 25$ periods. The first ten units are never treated while the last ten are subject to one treatment only. Other units can be treated once or twice.

The first treatment period is randomly drawn from a discrete uniform distribution with minimum 5 and maximum of 25 – meaning that no unit is treated in the first five periods – and it vary across units. The second treatment period, also randomly drawn from a uniform distribution, can be from one up to ten periods ahead of the first treatment. If, for a given unit, the second treatment occurs after the last sample period, the unit is treated only once.

Let x_{it} be the continuous treatment variable. For the first treatment, x_{it} is drawn from a uniform distribution ranging between 0.5 and 1. For the second treatment, x_{it} is drawn from a uniform distribution with minimum 0.3 and maximum 0.8.

The treatment effect is homogeneous across units, but heterogeneous over time. Given a change in x at time t, for a one unit increase in x the outcome variable increases by one at each horizon up to $t + 5$, when the treatment effect stabilizes (i.e., the treatment effect from horizon 0 to 5 is 1, 2, 3, 4, and 5). The second treatment has the same dynamic effect. Note however that, if the second treatment occurs before the first one stabilized (= before horizon 5), the second treatment effect is added on top of the still ongoing first treatment effect.

The outcome variable is therefore:

$$y_{it} = \alpha_i + \alpha_t + \sum_{l=0} \beta_l \Delta x_{it-l} + e_{it}$$

where α_i and α_t are unit and time effects, $e_{it} \sim N(0, 0.1)$ and $\beta_l = [1, 2, 3, 4, 5, ..., 5]$.

The results from 50 simulations are reported in Figure 14. Panel 14a plots the true average treatment effect together with the full range of treatment effects, together with the estimates of the clean LP-IV estimator, and its 95% multiplier. Panel 14b does the same but for the standard LP-IV estimator. We can see that the clean LP-IV tracks well the true average treatment effects at all horizons, and it is always within the full range of treatment effects. Differently, the standard LP-IV estimator underestimates the treatment effects at every horizon, and it is almost always outside the full range of results. This is confirmed by the Root Mean Square Error reported in panel 14c.

Figure 14: Monte Carlo Simulation Results

(a) Clean LP-IV Estimator (b) Standard LP-IV Estimator (c) Root Mean Square Error

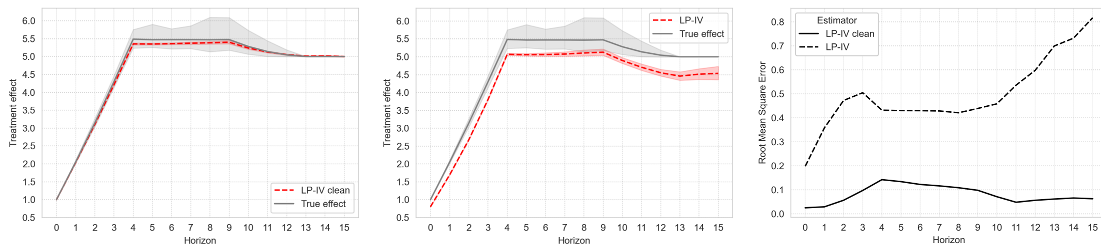

Results from the Monte Carlo simulations. Panel 14a plots the average true treatment effect (solid line) and the full range of results (shaded grey area), together with the treatment effects estimated with the clean LP-IV estimator and its 95% confidence interval. Panel 14b does the same but for the standard LP-IV estimator. Panel 14c reports the Root Mean Square Error.