The heterogeneous impact of the EU-Canada agreement with causal machine learning*

Lionel Fontagné † Francesca Micocci‡ Armando Rungi§

January 2025

Abstract

This paper introduces a causal machine learning approach to investigate the impact of the EU-Canada Comprehensive Economic Trade Agreement (CETA). We propose a matrix completion algorithm on French customs data to obtain multidimensional counterfactuals at the firm, product and destination levels. We find a small but significant positive impact on average at the product-level intensive margin. On the other hand, the extensive margin shows product churning due to the treaty beyond regular entry-exit dynamics: one product in eight that was not previously exported substitutes almost as many that are no longer exported. When we delve into the heterogeneity, we find that the effects of the treaty are higher for products at a comparative advantage. Focusing on multiproduct firms, we find that they adjust their portfolio in Canada by reallocating towards their first and most exported product due to increasing local market competition after trade liberalization. Finally, multidimensional counterfactuals allow us to evaluate the general equilibrium effect of the CETA. Specifically, we observe trade diversion, as exports to other destinations are re-directed to Canada.

Keywords: Free Trade Agreements; International Trade; Causal Inference; Machine Learning; Matrix Completion.
JEL Codes: F13; F17; C53; C55; L22

*We thank participants at the European Winter Meeting of the Econometric Society 2024. We are grateful for the useful comments by Gabor Bekes, Giovanna D'Inverno, Daniela Maggioni, and John Morrow. Lionel Fontagné thanks the support of the EUR grant ANR-17-EURE-0001. The views expressed in this paper are the authors'. Armando Rungi is grateful for the support of a PRIN grant 2022W2245M, funded by the European Union – Next Generation EU.

†Mail to: lionel.fontagne@univ-paris1.fr. University Paris 1, Paris School of Economics, 48 boulevard Jourdan 75014 Paris.

‡Mail to: francesca.micocci@imtlucca.it. Laboratory for the Analysis of Complex Economic Systems, IMT School for Advanced Studies, piazza San Francesco 19 - 55100 Lucca, Italy.

§**Corresponding author.** Mail to: armando.rungi@imtlucca.it. Laboratory for the Analysis of Complex Economic Systems, IMT School for Advanced Studies, piazza San Francesco 19 - 55100 Lucca, Italy.

1 Introduction

Ex-post estimates of the impact of Free Trade Agreements (FTAs) have been shown to be both unstable and fragile (Baier *et al.*, 2019). This can primarily be attributed to the challenges of effectively addressing issues of endogenous selection in trade agreements and the design of sensible counterfactuals. Due to the phasing-in of tariff reductions, staggered treatment adoption, where groups of products are treated over different periods, is an issue often raised when evaluating trade agreements (Nagengast & Yotov, 2024). And even if the design is not staggered, "forbidden comparisons" can be problematic if the treatment is not binary (De Chaisemartin & d'Haultfoeuille, 2023). These empirical challenges are all the more aggravated by the presence of heterogeneous firms in trade, which can sell multiple products and operate in multiple destinations.

In this contribution, we propose a causal machine-learning approach to uncover the impact of an FTA at the product and firm level. We apply this method to investigate the impact of the CETA (EU-Canada Comprehensive Economic and Trade Agreement) on French trade, using monthly customs data on the universe of French exports. Therefore, our empirical strategy evaluates multidimensional counterfactuals at the product, firm and destination levels. Following our proposed strategy, multidimensional counterfactuals are made possible by adapting a matrix completion algorithm for causal panel data originally suggested by Athey *et al.* (2021).

Notably, machine learning methods are increasingly used in economics for causal inference (Mullainathan & Spiess, 2017; Athey & Imbens, 2019). The simple intuition is that non-parametric methods are better at predicting potential outcomes in the presence of non-linearities by adopting less stringent assumptions on functional forms and the data-generating process. More specifically, we consider the French customs data as a matrix of observed outcomes to be partitioned between: i) treated vs. untreated observations, depending on whether the units of observation had seen a reduction of tariffs or a change in the quotas thanks to the CETA; ii) observations before and after the signature of the CETA.

Crucially, we can follow the application of the CETA agreement with monthly trade data from 2015M01 to 2018M12. As the signature occurred in September 2017, we split the timeline around that threshold. Then, we perform our exercise first at the product level, considering as treated the manufacturing products that have been included in CETA, and then at the firm level, this time considering multiproduct firms that have been concerned by the CETA because at least one of their products is enlisted by the treaty. In the product-level case, the matrix has cells identified by 5,118 products at the HS 6-digit level, 18 alternative destinations, including Canada, and time. In the second case, the matrix has each cell identified by 3,791 multiproduct firms, 18 destinations, up to three of their most exported products, and time.

Preliminary evidence suggests an endogenous product selection in the treaty, given that the products covered by the new CETA provisions already had, on average, a larger market for French producers before the treaty was signed. The products on which the parties negotiated were already exported by a greater number of firms in France, more frequently, with a lower average transaction value and a lower average value dispersion. We argue that such an endogenous selection needs to be monitored as it may be relevant for different trade policy environments. In our case, we implement a placebo test and confirm that matrix completion can handle endogenous selection.

Eventually, once the matrix of observed trade outcomes is designed, we can drop the observations of the treated units after the agreement entered into force and, thus, predict their trade values as if the CETA had not been signed. Crucially, predictions are obtained by exploiting all the information left in the matrix, including two years before the treaty. On the other hand, we can control the prediction accuracy of the method by looking at the elements of the matrix that were not treated. Following standard approaches in machine learning methods, we train our model on five random folds of the part of the matrix that includes untreated units, and then we check out-of-sample how far our predictions are from actual realizations of the outcomes.

For our purposes, CETA is a compelling case of an FTA whose negotiation has been intricate, lengthy and contrasted. It took ten years since the first discussions[1] to have the agreement provisionally entered into force in 2017. According to its provisional enforcement, most of the trade provisions in the agreement have already been applied, although it is still awaiting final ratification by all EU members.[2] During the negotiations, France emerged as one of the main proponents of establishing a closer trading relationship with Canada. A shared colonial past, a common language,[3] and similar consumer preferences give Canada more than an incentive to trade with France. Ratification by the French Assembly was voted on in July 2019, and the agreement was examined and eventually rejected by the Senate in March 2024.

Yet, an asymmetry was evident from the beginning for all parties involved in the negotiation. The treaty would have *prima facie* been more relevant for Canada than for European countries. However, the EU's interest was to foster unprecedented economic cooperation

[1] It dates back to a Canada-EU bilateral summit in Berlin in 2007.

[2] Even if the European Commission is solely in the competence of the trade policy of the European Union, in July 2016 it was decided that CETA qualified as a *mixed agreement* because it touches upon other policy domains different from trade, and thus it needed to be ratified also through national procedures.

[3] English and French have been established as joint official languages since 1969.

with new partners in the face of the rise of emerging markets like China (Hübner *et al.*, 2017) and to have a testing ground for *deep trade agreements* covering areas beyond tariffs. Notably, an asymmetry in the size of parties involved in the Treaty makes the local competition among European exporters potentially larger compared to the relatively smaller positive demand shock induced by the trade liberalization.[4] Therefore, by looking from the perspective of a single exporting country, France, we would expect a non-negligible impact, possibly magnified by the competition of French exporters with other European producers.

We proceed with our investigation in three steps. At first, we evaluate the overall impact of CETA at the product level. Crucially, at this stage, we find that CETA positively impacted French exports at both the intensive and extensive margins. On the one hand, product-level flows to Canada increased on average by 1.28%. On the other hand, we find that there has been a relevant product churning due to the treaty beyond regular entry-exit dynamics, as about 13.1% of new French products reached Canada for the first time, and 11.9% of them abandoned the market thanks to the new provisions.

Importantly, our matrix completion approach allows us to expose the relevant heterogeneity of the impact of a trade treaty. We argue that it is an advantage with respect to other more synthetic empirical strategies. In fact, we can evaluate the full distribution of treatment effects that emerge from the matrices, i.e., on each product or firm that is concerned by the CETA. In doing so, we observe that we have both cases of positive and negative impacts on observed units and that, for example, the treatment effects on single products are positively associated with a measure of revealed comparative advantage for French exporters vs. the rest of the world. That is, the increase in the export flow has been higher for those products for which French producers had a competitive edge before the treaty signature. Similarly, when we consider the heterogeneity at the extensive margin, we find that product churning is positively associated with the elasticity of substitution. In other words, as largely expected, the French products that either enter or exit the Canadian market as a direct consequence of the new treaty are also the ones that have an elasticity of substitution that is higher if compared with products that just continue to be exported.

In the second stage of the analysis, we investigate the firm-level dimension with a special focus on the strategies of multi-product firms. Trade theory tells us that the latter can adjust their portfolios after the signature of a trade treaty. After we rank products within firm-level portfolios, we find that multi-product firms, on average, sell relatively more of the already first-sold products to Canada after the CETA. We believe this result is consistent with the

[4]Please note that Canada's GDP is similar in size to Italy's. France is Canada's ninth-largest trading partner and the fourth-largest among EU members. At the same time, Canada stands as only the thirtieth-largest partner, amounting to a share of only 0.8% total exports.

theoretical framework proposed by Mayer *et al.* (2021) and Eckel & Neary (2010), according to which multiproduct exporters tend to reallocate their product mix as a response to the demand shock in the export markets. In fact, trade liberalization generates relatively higher competition for French exporters, who find it convenient to invest relatively more and focus on the products on which they have a higher competitive advantage.

Finally, we follow best practices in the trade literature dealing with general equilibrium effects of a change in bilateral trade costs between parties to a trade agreement (Head & Mayer, 2014; Anderson & Yotov, 2016). Cancellation of tariffs between the parties increases relative trade costs between the parties and third countries, leading to indirect trade effects. This is indeed consistent with a classical Vinerian diversion effect (Viner, 1950), whereby trade between parties to a PTA partially substitutes for trade between third parties that do not participate in the PTA. Following this logic, reducing trade costs with Canada is equivalent to a relative increase in the costs of exporting to other destinations. In our context, trade diversion takes the form of indirect policy spillovers: we detect a significant and negative association between the effects on the export of products from France to Canada enlisted by the CETA and the changes in the exports of the same products from France to alternative destinations. The correlation is all the more significant for products with a relatively higher substitution elasticity.

The remainder of the paper is structured as follows. We begin with a short review of the relevant literature in Section 2. Section 3 presents the data and offers preliminary evidence. In Section 4, we outline the empirical strategy. Results are displayed in Section 5, while robustness and sensitivity checks are presented in Section 6. Section 7 concludes.

2 Related Literature

Ex-post evaluation of free trade agreements is challenging (Baier *et al.*, 2019) because they often entail an endogenous selection of partners or products (Baier & Bergstrand, 2004, 2009), on the one hand, and a self-selection of heterogeneous exporters (Melitz, 2003), on the other hand. Hence, Goldberg & Pavcnik (2016) consider this endogeneity a major hurdle to the causal identification of the economic impact of FTAS.

This endogeneity of PTAs has been addressed by using various approaches, including gravity equations with additional controls (e.g. bilateral fixed-effects) for unobserved characteristics (Aitken, 1973; Abrams, 1980; Bergstrand, 1985; Soloaga & Wintersb, 2001; Feenstra *et al.*, 2001), instrumental variable (IV) or control-function techniques with cross-sectional data (Baier & Bergstrand, 2002; Magee, 2003; Baier & Bergstrand, 2009), panel data models with a rich set of fixed effects (Head & Ries, 1998; Baier & Bergstrand, 2007; Westerlund

& Wilhelmsson, 2011; Yang & Martinez-Zarzoso, 2014), or matching techniques (Baier & Bergstrand, 2009; Egger & Tarlea, 2021).[5]

In this paper, we explore the scope for using a potential outcome model to assess the causal impact of preferential trade agreements.[6] In particular, we draw from the most recent advances in causal machine learning, whose aim is to estimate average causal effects after predicting the missing potential outcomes with non-parametric methods (Abadie *et al.*, 2010, 2015; Arkhangelsky *et al.*, 2019; Chernozhukov *et al.*, 2021). Specifically, we leverage the literature on matrix completion that originally exploited observed information to predict unobserved information when matrices are sparse (Candes & Plan, 2010; Candes & Recht, 2012; Mazumder *et al.*, 2010). For our purpose, we adapt the algorithm initially proposed by Athey *et al.* (2021), whose intuition is that a matrix approach can also be used for causal inference while allowing for time dependence, unregularized units and time-fixed effects. All properties that, according to Athey *et al.* (2021), help boost the quality of potential outcomes' predictions.

On top of empirical challenges, we know from trade theory that opposing mechanisms may hinder an accurate estimate of the impact of tariff reduction. On the one hand, a tariff reduction implies greater market access because the demand increases in the liberalized market. On the other hand, tariff reductions under trade agreements may have pro-competitive effects. When Marshall's second law of demand does not apply, monopolistic exporters may reduce their markups in response to reduced tariffs (Mrázová & Neary, 2017) or preferential market access (Crowley & Han, 2022). This induces, in turn, selection effects. Market size and trade openness affect the intensity of competition in a market, which reinforces the selection of exporters to that market (Melitz & Ottaviano, 2008). Against this background, we design our empirical strategy encompassing multidimensional counterfactuals, both at the product and firm level, which enable us to discuss competing mechanisms.

Crucially, our empirical design acknowledges the role of heterogeneous firms in trade agreements, especially in a world where multi-product firms dominate trade flows (Feenstra & Ma, 2007; Eckel & Neary, 2010; Iacovone & Javorcik, 2010; Bas & Bombarda, 2013). In this, we refer to Mayer *et al.* (2014) and Bernard *et al.* (2010, 2011), who incorporate multi-product firms into models of heterogeneous firms while building upon the pioneering work by Melitz (2003). They show that tougher competition in a liberalized market leads

[5]See the reviews by Limão (2016) and Larch & Yotov (2023) of the empirical exercises estimating the impact of trade agreements.

[6]The framework for causal inference that uses 'potential outcomes' to define causal effects at the unit level in the context of randomized experiments and quasi-experiments is dubbed Rubin Causal Model (Rubin, 2005). The introduction of this framework in economics helped comply with the so-called *credibility revolution* cited by Angrist & Pischke (2010).

firms to skew their export sales towards their better-performing products. On a similar line of research, Dhingra (2013) and Qiu & Zhou (2013) predict that falling trade costs make the most productive firms expand their product scope, and the least productive firms contract theirs. According to Baldwin & Gu (2009), the net effect could be ambiguous because tariff cuts can both increase exporters' plant size by extending the production-run length of the exported portion of the product line and reduce the exporters' plant size by reducing the total number of products.

A final layer of complexity that we consider in this contribution arises when considering the adjustment mechanisms of firms to multiple destinations. Two mechanisms concur with third-country effects, i.e., on destination markets that are not part of the signed PTA. On the one hand, reducing trade costs between the EU and Canada increases the relative cost of exporting to countries that are not parties to the agreement. General equilibrium effects of a change in the matrix of bilateral trade costs are conducive to indirect trade effects (Head & Mayer, 2014; Anderson & Yotov, 2016). Trade between parties to a PTA partially substitutes for trade between parties and third countries, which should appear at the aggregate level (Viner, 1950). On the other hand, at the firm level, the determinants of exporters' geographical expansion reveal patterns of entry, sales distribution across markets, and export participation (Eaton *et al.*, 2004, 2011, 2012; Eaton & Fieler, 2019). Notably, Arkolakis & Muendler (2013) and Arkolakis *et al.* (2021) found that the scope of exporters is unrelated to the size of destination markets, but it is related to geographic distance. As a result, after trade liberalization, we expect to observe a larger effect on the intensive rather than the extensive margin of trade depending on the geographical distance of the trading partner.

3 Data and preliminary evidence

3.1 Customs data and trade regime changes

Our primary data source is the French Customs (*Direction Générale des Douanes et Droits Indirects*)[7], where we have records of trade values at the product, firm, and month levels. Products are originally classified by the 8-digit Combined Nomenclature (CN8), and firms are identified by their *SIREN* number, i.e., the 9-digit identifier assigned to every registered business in France by the National Institute of Statistics and Economic Studies. Moreover, we rely on the WTO tariff databases to retrieve information on those products at the HS 6-digit level whose tariffs or tariff quotas have been modified by the EU-Canada Comprehensive

[7] The database was accessed through the CASD, French Secure Data Access Center (project DYNAMEX).

Economic and Trade Agreement (CETA).[8]

Original customs data are first aggregated from monthly to yearly levels in September-August segments, following the timeline of the trade treaty, which became operational in September 2017. In addition, we align the product classification from the 8-digit Combined Nomenclature (CN) to the 6-digit Harmonized System (HS) classification to match the original information on products whose tariff or tariff quota has been changed by CETA. Since the HS classification was revised in 2017, we converted the codes of entries back to HS 2012.

So far, we have identified the perimeter of the product-level analyses we perform in Section 5.1. Our investigation encompasses all products that France exports to Canada regardless of the firms' characteristics. In the second part of the empirical strategy, we will focus on the impact that CETA has on multiproduct firms; therefore, we need to eliminate from our sample perimeter[9]: i) firms that do not export to Canada, ii) firms that export only one product to Canada.

In Figures 1 and 2, we provide waterfall charts to visualize the relevance of products and firms included in our study. On the one hand, when we separate products liberalized after CETA, we observe that they make up 77% of the total product lines exported from France to Canada. On the other hand, the list of products that have seen a change in the tariff or non-tariff regime thanks to CETA coincides for about 57% with the list of product lines that French exporters already trade with the rest of the world.

[8] Appendix Table A1 briefly summarizes the extent of tariff changes for French exporters in Canada due to CETA.

[9] In the original data, we find firms that are active in service industries and occasionally export goods. We eliminate these cases from our firm-level sample perimeter because they conceal a delivery of materials needed to proceed with the service supply (e.g., building materials for construction firms, laboratory equipment for an R&D company, etc.).

Figure 1: Products' coverage in 2016

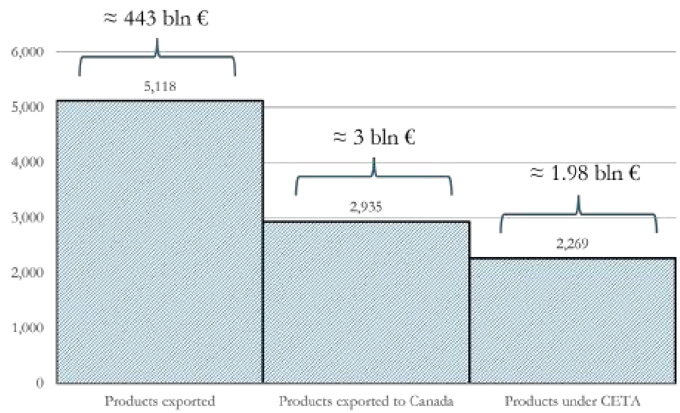

Note: The figure shows sample coverage of products in 2016. The y-axis indicates the number of products, whereas the text boxes on top of the bars indicate the total trade value in 2016. On the left is the number of products exported from France to any destination. In the centre is the number of products exported to Canada. On the right is the number of products that are both exported to Canada and fall under the provisions of the Canada-EU Trade Agreement.

Figure 2: Firms' coverage in 2016

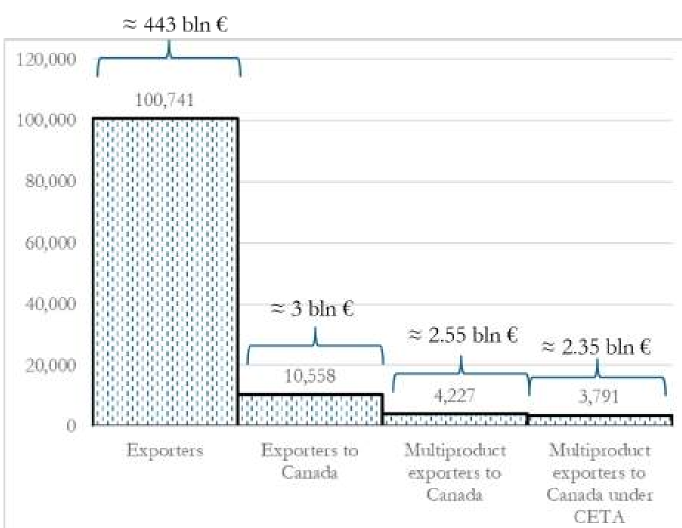

Note: The figure shows sample coverage of exporters in 2016, while text boxes on top of the bars indicate the total trade value in 2016. On the left is the number of French firms that exported to any destination. Then, we report the number of exporters to Canada and, among the latter, the number of multiproduct firms because they export at least two products to Canada. On the right is the number of multiproduct exporters to Canada, with at least one product enlisted by the Canada-EU Trade Agreement, for which we indicate the value of their total exports to Canada, encompassing both products with and without a trade regime change.

From our perspective, either stylized fact is worth further investigation. In the first case, we expect an endogenous selection of products in the treaty negotiation, and we test it in the following paragraphs. In the second case, we expect general equilibrium effects inducing indirect trade effects on alternative destinations as it will be *relatively* more costly to export the same products to alternative destinations: see Section 5.3.

As for firms, we first need to drop those that have never exported to Canada because they are not directly concerned about the signature of the CETA. Then, following a basic definition of multiproduct firms, we will consider only those that export at least two products to Canada. In this case, as from Figure 2, we can see that only about 10.5% of French exporters reach Canada as an export destination. Among them, about 40% are multiproduct firms and can sell a portfolio of at least two products in Canada. Finally, among the latter, 79.8% have seen a tariff or non-tariff change in at least one of their products exported to Canada after CETA.

In the second part of the paper, the subset of multiproduct firms is of special interest to us not only because they are relevant in terms of aggregate trade flows (2.55 billion euros vs 3 billion euros of total exports to Canada) but also because they are a segment that potentially shows adjustments in product scope, which would be otherwise hidden if we do not consider the firm-level dimension. In Appendix Figure A1, we show French exporters' distribution of product portfolios to Canada.

3.2 Preliminary evidence

In the following paragraphs, we investigate whether products and firms that have seen a change in the trade regime significantly differ from those that have not. The obvious intuition is that negotiators could have picked production segments that could show higher gains from trade. Alternatively, it is possible that bigger firms had the power to impose their own agenda on negotiators. In Table 1, we investigate the issue with two sequences of t-tests on the difference in means of indicators that could possibly capture the peculiar differences between products included and not included in the CETA. First, we test our indicators considering bilateral exports from France to Canada. Then, we consider the same partition of products under the CETA, this time looking at the features of products and producers at the global level after aggregating over destinations.

The first three indicators we test in Table 1 refer to features of the product-level monthly flows observed in the period 2015M01 to 2016M12, while the other two indicators refer to the firm-level dimension. Starting from the top of the table, we observe that the average trade value of products included in the CETA had a lower magnitude, a lower dispersion around the sample means, and its transactions were more frequent in the two years preceding the treaty's signature. If we look at exporters, the product was usually traded by more firms,

Table 1: Characteristics of trade flows before CETA - 2015M01-2016M12

	products in the CETA	products not in the CETA	difference in means
Exports to Canada			
Avg. trade value	30231.8	54023.6	-35700.5***
Avg. dispersion	65579.8	122571.7	-78671.4***
Avg. number of transactions	2571.4	599.9	1971.4***
Avg. number of firms	212.1	100.2	111.8***
Avg. firm's exports	509,037.5	207,466.9	301,507.6***
All exports			
Avg. trade value	35265.7	60645.2	-25379.5***
Avg. dispersion	162147.3	301385.0	-188687.6***
Avg. number of transactions	42852.1	23216.9	19635.2***
Avg. number of firms	1290.5	1278.3	12.18***
Avg. firm's exports	8,150,142	1,412,479	6,737,762***

Note: The table reports t-tests computed on average indicators of the export matrix in 2015-2016 considering products that will see a change with CETA in 2017 (column 2) vs. products whose trade regime will not change (column 3). Column 4 reports differences in the means considering unequal variances. *** stands for $p \leq 0.001$, hence the average means are significantly different. In the first half of the table, we consider only export flows to Canada, i.e., the destination involved in the treaty. In the second bottom half of the table, we enlarge the matrix to consider export flows to all export destinations, although they are not parties in the CETA.

which had, on average, a relatively higher exposure to Canada as an export destination.

If we look at the bottom of the table, we see that the same differences observed in the bilateral relationship between France and Canada are confirmed by aggregate flows between France and the rest of the world. Briefly, products included in the CETA are usually traded by firms whose export size is, on average, bigger, while single monthly flows are smaller, more frequent, and with lower volatility around the mean value in the two years before the CETA.

Eventually, preliminary evidence shown in Figure 1 motivates the choice of an empirical strategy that is capable of handling an endogenous selection of product lines in a trade treaty, thus making policy evaluation unbiased by the political economy of the bigger firms or by the tendency of negotiators to cherry-picking products that already have a higher potential.

Our preferred empirical strategy should also be capable of handling the presence of heterogeneous time trends. It is, in fact, possible that products and firms concerned by the CETA were already on paths to growth before the treaty was signed. The presence of unparallel time trends could possibly confound the actual impact of the trade treaty. In Figures 3 and 4, we display linear trends after the estimation of simple difference-in-difference mod-

Figure 3: Time trends at the product level, intensive and extensive margins

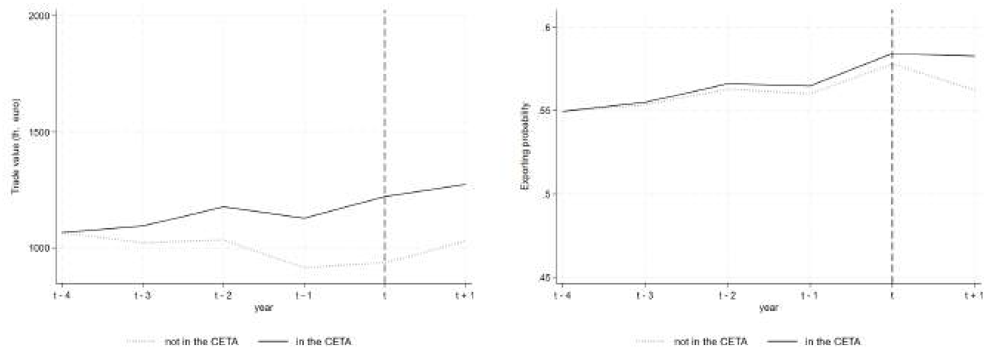

(a) Intensive margin (b) Extensive margin

Note: We report in panel (a) linear trends for trade values of product lines exported to Canada, separating those that are included in the CETA and those that are not. In panel (b), we report linear trends for the probability that a new product line is exported to Canada, separating those that are included in the CETA and those that are not. The graphs are generated using the predictions of a difference-in-difference model augmented with interactions of time with an indicator of treatment when products are enlisted by the CETA.

Figure 4: Time trends at the firm-level on the intensive and extensive margins

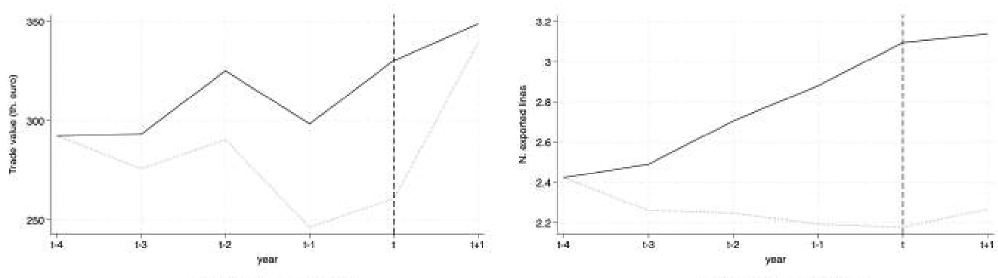

(a) Intensive margin (b) Extensive margin

Note: We report in panel (a) linear trends for trade values of firms that exported to Canada, separating those that have a product enlisted by the CETA and those that have not. In panel (b), we report linear trends for the number of lines a firm exports to Canada, separating those that have a product included in the CETA and those that have not. The graphs are generated using the predictions of a difference-in-difference model augmented with interactions of time with an indicator of treatment when firms have a product enlisted by the CETA.

els[10] of the intensive and extensive margins for both products and firms, separating when they are concerned by the CETA and when they are not.

After a graphical inspection, we can observe that intensive margins at the product and firm levels (panels (a) in Figures 3 and 4) were already on diverging paths. In the case of products, those not included in the CETA were already on a downward trend. In the case of firms, those that do not have a product enlisted by the CETA had been on a decreasing trend in the years before the treaty and then increased significantly thereafter. In the case of extensive margins, product flows do not show significant differences, while firm-level pre-trends were significantly diverging.

4 Empirical strategy

4.1 Treated products and treated firms

In the following paragraphs, we develop an empirical strategy to evaluate the impact of CETA. For the sake of generalization, we will define a generic u-th unit of observation at time t, such that the exposure to CETA, i.e., our treatment, can be defined as W_{ut}. Yet, for our purpose, we need to introduce two different definitions of policy treatment: one at the product level and one at the firm level.

At the product level, we will consider the treated population, $\mathcal{T}$, consisting of all the products that experienced a tariff or a quota change after CETA. Let p denote the product, d represent the destination, and t indicate time. Notice that d can indicate either Canada, as it is the only destination in which treated products are exported with a tariff or quota change, or it can indicate alternative destinations different from Canada. Please note that we consider a product as treated regardless of the destinations in which it is exported. The latter setup will turn out to be useful when we evaluate general equilibrium effects later in the paper.

Since CETA entered into force in September 2017, we aggregate monthly flows by year τ

[10] We estimate simple difference-in-difference models augmented with terms that capture the differences in slopes across the products/firms that are concerned by the CETA and those that are not. See Appendix B for more details. Results of the difference-in-difference models are reported in Appendix Table B1. Please note how diff-in-diff results suggest that the CETA had only an effect on the firm-level extensive margin, whereas no significant impact is registered on the intensive margins at the product and firm levels. While serving as a valuable reference point, a simple diff-in-diff methodology cannot be valid if the assumption of parallel trends is violated, as from Figures 3 and 4, and when the treatment is not orthogonal to relevant characteristics of the treated units, as from Table 1.

in the period September-August.[11] In this case, the treatment indicator is defined as follows:

$$W_{pdt} = \begin{cases} 1 & p \in \mathcal{T}, t \geq \tau \\ 0 & otherwise \end{cases} \tag{1}$$

When we switch to the firm level, our population consists of multi-product firms that export to Canada at least two distinct products.[12] Among them, the set of treated firms Θ is defined as:

$$\Theta = \{i: \ \Psi_{itCA} \cap \mathcal{T} \neq \emptyset, t = [\tau - 2, \tau], \}$$

where Ψ_{itCA} represents the set of products p exported to Canada by firm i in year t, and $|\Psi_{itCA}| \geq 2$. Briefly, we consider as treated any firm that, before or after the entry into force of CETA, exported at least two products to Canada[13], with at least one of them enlisted by the CETA. Conversely, we will consider non-treated firms that exported at least two products to Canada before CETA but do not have in their portfolio any products included in the CETA.

Once we have defined the set of *treated firms*, Θ, we can establish the treatment at the firm-*per*-product level. Let i denote the firm, p indicate the product, and t represent the year. The treatment indicator at the firm level is defined as:

$$W_{ipt} = \begin{cases} 1 & \forall \ |\Psi_{itCA}| \geq 2, \ i \in \Theta, \ t \geq \tau \\ 0 & otherwise \end{cases} \tag{2}$$

.

Therefore, in the following paragraphs, when we deem it not necessary to specify it, our generic indicator of treatment W_{ut} for the u-th unit will suffice. When presenting results, we will indicate which of the eqs. 1 or 2 defines the treatment.

4.2 Matrix completion

At this point, we are ready to illustrate the details of our causal machine-learning application on trade policy evaluation. Originally, matrix completion methods were used to recover lost

[11] In the following, $\tau - 2$ refers to the period from September 2015 to August 2016, $\tau - 1$ refers to the period from September 2016 to August 2017, and τ refers to the period from September 2017 to August 2018. Our dataset provides information up to December 2018, which means we can only observe one period (τ) ahead of CETA. Consequently, the analysis is restricted to the short-term effects of the Treaty. Nonetheless, our approach is also suitable for analyzing a staggered adoption scheme across multiple post-treatment periods.

[12] See Section 3 for a description of the firm-level sample selection strategy.

[13] Note that in the following, CA stands for Canada.

information in highly sparse matrices. In the context of statistical and computer science exercises, the task has been to fill in the missing entries of a matrix that was only partially observed (Candes & Plan, 2010; Mazumder *et al.*, 2010; Candes & Recht, 2012). The novel intuition by Athey *et al.* (2021) is that one could instead frame a matrix completion algorithm in the context of potential outcome models with predictions of missing multidimensional counterfactuals. We adapt the framework by Athey *et al.* (2021) to our case of trade policy evaluation, when we have N units of observations (products or firms), T time periods, and there exists a pair of potential outcomes, $Y_{ut}(0)$ and $Y_{ut}(1)$, with unit u exposed in period t to the entry into force of the CETA. The generic treatment has been defined in the previous section as a matrix with entries $W_{ut} \in \{0, 1\}$, and the realized outcomes are thus equal to $Y_{ut} = Y_{ut}(W_{ut})$.

In our case, the fundamental problem of causal inference is that a set $\mathcal{M} < NT$ of potential outcomes is not observed. Specifically, we do not observe the outcomes of the treated units as if the treatment did not occur. In our context, we will never observe the potential exports of products or firms concerned by the CETA as if the latter was not signed. Briefly, we need valid counterfactuals for the set $\mathcal{M}$, and the solution is to predict them using the information available in the trade matrix from entries $\mathcal{O} \equiv NT - \mathcal{M}$, which are observed. Once we obtain valid counterfactuals, we can compute the relative treatment effect on the treated (TET) expressed in monetary values as:

$$\forall\{u,t\} \in \mathcal{M} : TET_{ut} = Y_{ut}(1) - \hat{Y}_{ut}(1) \tag{3}$$

.

Then, we can manipulate the latter expression to find the best solution, in levels or in percentage points, depending on whether we want to comment on the intensive or extensive margin, as we explain in the following paragraphs.

4.2.1 Effects on the intensive margin

We can evaluate the impact of the new trade regime on the intensive margin after looking at the moments of the entire distribution produced by the entries we obtain from the matrix of counterfactuals. In this case, we prefer to express the treatment effect on treated from eq. 3 as a ratio, to comment in relative terms and on percentage points, in the form:

$$\forall\{u,t\} \in \mathcal{M} : TET^{*}_{ut} = \frac{Y_{ut}(1) - \hat{Y}_{ut}(1)}{Y_{u,t-1}(1)} \times 100 \tag{4}$$

where Y_{ut} is the observed value for unit u at time t, $\hat{Y}_{ut}$ is corresponding predicted value, and $Y_{u,t-1}$ is the observed value for unit u at time $t-1$ Finally, we can compute the weighted average treatment effect on the treated (WATET), also expressed in relative terms, in the form:

$$WATET = \sum_{\{u,t\} \in \mathcal{M}} s_{ut} TET^*_{ut} \tag{5}$$

,

where s_{ut} indicates the salience of the export flows. For the sake of simplicity, we can use for salience the share of the trade flows of unit u at time $t-1$, i.e., before the signature of the CETA, on the total export flows for each entry $\{u, t\} \in \mathcal{M}$.

4.2.2 Effects on the extensive margin

In the evaluation of the extensive margin of trade, the potential outcomes are binary, $Y_{ut}(1) = \{0, 1\}$, i.e., they are equal to one if the product is exported and zero otherwise. Our matrix completion application reduces to a classification problem, and we obtain predictions in a binary form, $\hat{Y}_{ut}(1) = \{0, 1\}$, such that treatment effects can have three alternative values, $TET_{ut} \in \{-1, 0, 1\}$. A value -1 means that our counterfactual predicts that a trade flow existed in that entry of the trade matrix, but it actually did not. We will define the latter as the negative extensive margin. A value of 1 implies that our counterfactual indicates that the product should not have been traded, but it actually was. We will call the latter the positive extensive margin. On the other hand, every time that we find a $TET_{ut} = 0$, it means that our counterfactuals and the observed outcomes corresponded. Please note that, against the previous background, products can still enter or exit the foreign market following regular product churning, regardless of a change in the trade regime. The latter cases would all be flagged with a zero in the set of treatment effects.

4.2.3 The estimator

Let us start by representing the entire trade matrix from the original data. In the product-level analysis, we will have a matrix with entries defined by the trade value of each 6-digit product-*per*-destination (i.e., the u-th observation) and time in a cell. In the firm-level analysis, we report each matrix cell's trade by firm-*per*-product (i.e., the u-th observation) and time. Next, we empty the set $\mathcal{M}$ of matrix entries where we have exports with tariff and tariff-quota changes after the CETA signature, i.e., $Y_{ut}(1)$ when $\geq$ 2017, and we ask the algorithm to reconstruct the full matrix while feeding it information from the set $\mathcal{O}$, including:

1. treated and untreated observations before the treatment, when CETA did not exist (i.e., $Y_{ut}(1)$ and $Y_{ut}(0)$ when $t < 2017$)

2. untreated observations after the treatment (i.e., $Y_{ut}(0)$ when ≥ 2017)

Further details on the product-level and firm-level trade matrices are described in Sections 5.1 and 5.2, respectively. In our context, the value of a matrix completion approach lies in its ability to leverage non-parametrically all available information without making stringent assumptions on joint distributions and functional forms. By predicting each unobserved potential outcome, we obtain multidimensional counterfactuals for each cell in a matrix that pertains to treated units, therefore taking on board all the heterogeneity that can possibly derive from a trade policy treatment.

We obtain predictions from a decomposition of the $N \times T$ matrix $\mathbf{Y}$, such that:

$$\mathbf{Y} = \tilde{\mathbf{Y}} + \tilde{\gamma} + \tilde{\delta} + \varepsilon \tag{6}$$

where we can collect $\hat{\mathbf{Y}} = \tilde{\mathbf{Y}} + \tilde{\gamma} + \tilde{\delta}$, as these are the components we want to estimate. Among them, $\tilde{\mathbf{Y}}$ is a low-rank matrix with respect to the original $N \times T$. Then, we have $\tilde{\gamma}$, which is the $N \times 1$ vector of row-fixed effects, and $\tilde{\delta}$, which is the $1 \times T$ vector of time fixed effects.[14] In our context, the $N \times 1$ vector of row-fixed effects can represent either product-destination or firm-level fixed effects, respectively. Eventually, we leave ε as an $N \times T$ matrix of random noise values.

Our $\tilde{\mathbf{Y}}$ is the result of a singular value decomposition (SVD), such that $\tilde{\mathbf{Y}} = \mathbf{S}\boldsymbol{\Sigma}\mathbf{R}^{\top}$, where $\mathbf{S}$ and $\mathbf{R}$ are unitary matrices, and Σ is a rectangular diagonal matrix with singular value entries $\sigma_u(Y)$. The latter entries are substituted by $\max(\sigma_i(\tilde{\mathbf{Y}}) - \lambda_Y, 0)$ after regularization. In fact, we introduce regularization on the $\tilde{\mathbf{Y}}$ component, $\lambda_Y||\tilde{\mathbf{Y}}||$, to avoid overfitting. In our context, overfitting would imply that the model corresponded too closely to the training matrix, and its power would be poor in predicting counterfactuals. Indeed, overfitting problems more likely arise in cases like ours where we have a high $N \times T$ dimensionality. Finally, the estimator can be written as the result of an optimization problem in the general form:

$$\min_{\tilde{Y},\gamma,\delta} \left[\sum_{(u,t)\in\mathcal{O}} \frac{1}{|\mathcal{O}|} \left(Y_{ut} - \tilde{Y}_{ut} - \gamma_u - \delta_t \right)^2 + \lambda_Y ||\tilde{\mathbf{Y}}||_* \right] \tag{7}$$

[14]Note that the row and column-fixed effect can be subsumed in matrix $\tilde{\mathbf{Y}}$. However, Athey *et al.* (2021) already pointed out that separating fixed effects without regularization greatly improves prediction quality. In our case, we confirm that prediction power deteriorates when we do not separate fixed effects.

where $\mathcal{O}$ includes any pair (u, t) in the set of observed export outcomes, and $||\tilde{\mathbf{Y}}||_*$ is the nuclear norm of the matrix $\tilde{\Sigma}$ resulting from shrinking the scaling matrix with the singular value decomposition (SVD) by λ_Y. We select the optimal value of λ_Y after cross-validation[15] on K different random subsets $\mathcal{O}_k \subset \mathcal{O}$ of the original matrix, having a fraction of observed data equal to the one in the original sample. Finally, once we have predicted matrix $\hat{\mathbf{Y}}$, we obtain the counterfactuals we need to estimate treatment effects as in eq. 3.

5 Results

In this section, we discuss the findings of our application to both a product-level and a multiproduct firm-level investigation. For each case, we introduce separate exercises for the intensive and extensive trade margins. In each case, we start by describing the specific design of the matrix structure that we draw before running the estimator. Then, we report the prediction accuracies always needed to validate the model. Finally, we comment on the results with the help of a few post-estimation statistics.

5.1 Product-level analysis

The unit of observation is the product p at the 6-digit level of the HS classification exported at time τ to different destinations d. A product is *treated* if its tariff or quotas have changed after CETA since September 2017[16]. Therefore, in this section, we are interested in evaluating treatment effects on the treated in percentage points, which we now write as TET^*_{pdt} because the general u-th unit of observation is now represented by a product p, at destination d, and time t.

For our purpose, besides Canada, we aggregate and rank major destinations of French exports to avoid matrix sparsity.[17] We compute two separate destination rankings, and then we consolidate them. At first, we rank importing countries based on the average total trade value they received from France in 2010-2016. In a second exercise, we rank destinations after counting the number of products received from France in the same period. Finally, we include in our selection those countries that are in the top ten in either ranking. The

[15] As we choose a nuclear norm for regularization, the estimator can be computed using fast convex optimization programs like the one proposed by Mazumder *et al.* (2010).

[16] Please note how, since eq. 1, we consider the treatment to be product-specific and not destination-specific. The reason is that we will also investigate policy spillovers in destinations that are not directly affected by the CETA, as it will become evident in Section 5.3.

[17] As in Fontagné *et al.* (2018), we also observe a high sparsity because the selection of products at each destination is stringent. In the original data, the vector of products exported to each destination contains, on average, at least 80% of zeros. A highly sparse matrix with an inflation of zeros complicates calculations while saturating computer memory.

Figure 5: Matrix Structure for the product-level analysis

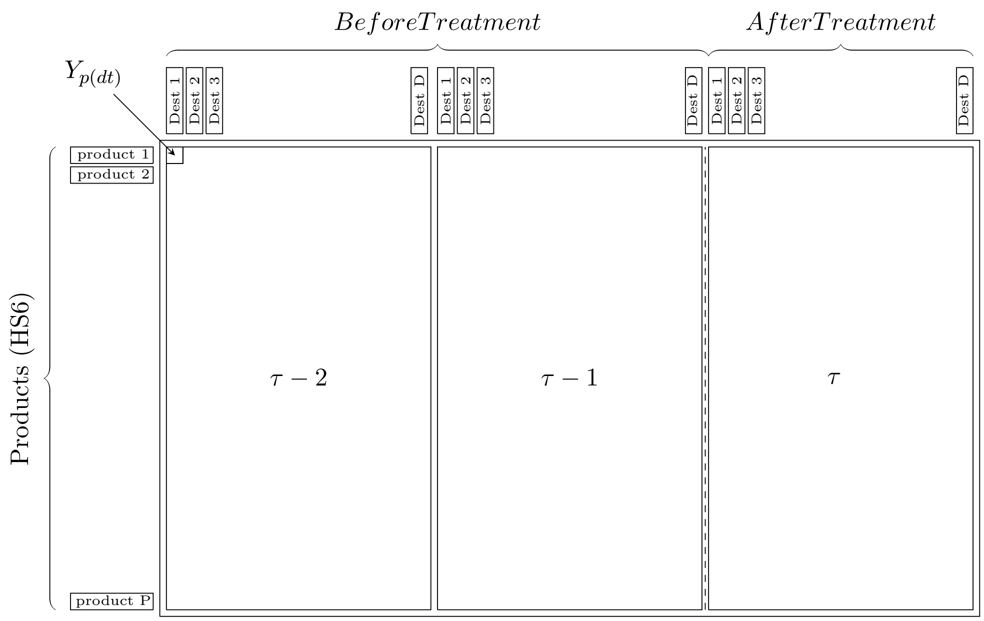

remaining destinations are mainly aggregated by continent (e.g., the rest of Europe, the rest of Asia, etc.). In Appendix Table A3, we record the relative trade importance of each destination in our final ranking.

As for products, we ensure we can properly separate the intensive and the extensive margin. In the first case, we only consider the subset of products that were exported to Canada in either of the two years before the treatment and were still exported after the CETA.[18]

In Figure 5, we visualize our matrix structure. In the case of the intensive margin, the P rows of the matrix correspond to the HS 6-digit products exported by France. The TD columns of the matrix, instead, correspond to the set of D possible export destinations in T different times. Then, each matrix element Y_{pdt} is the total export value for product p at destination d and time t.

In the case of the extensive margin, our focus is the effect on the export probability of treated products. In this case, we will consider all possible products $\mathcal{P}$ exported by France anywhere, and each matrix element is a binary variable, $Y_{pdt} = \{0, 1\}$, which takes the value 1 if product $p \in \mathcal{P}$ is exported at destination d in time t, and 0 otherwise.

We estimate the model by solving the minimization problem described in the generic eq. 7, and we obtain two matrices of predicted outcomes: one for the intensive margin and one for the extensive margin. Then, crucially, Table 2 reports some measures of the prediction

[18] For a visual representation of the trade patterns included in the intensive margin, see Appendix Table A2.

Table 2: Prediction accuracy at the product level

model	min RMSE	$\overline{Y}$	SI	NRMSE
Intensive Margin	7.12126	7,060.71	0.000001	0.00027
Extensive Margin	0.25861			0.25861

Note: The table reports standard measures of prediction accuracy. $\overline{Y}$ is the average trade of a line p in a year for any destination d, and it is used to compute the normalised version of the RMSE and the Scatter Index (SI). The value of $\overline{Y}$ indicates the average predicted counterfactual in monetary values. On the extensive margin, no normalization is required, as the predicted outcomes are already in a range $0, 1$.

accuracy. Briefly, a certain level of prediction accuracy guarantees that our empirical model returns valid counterfactuals. If the predicted values are close enough to the observed values, then we expect a minimum bias when we evaluate the impact of the policy. As in a standard machine learning framework, the algorithm is first trained on different in-sample subsets and then evaluated on out-of-sample segments. In our specific case, the evaluation is made with a minimum average Root Mean Squared Error (RMSE) obtained after five random folds.[19]

Notably, we record a high prediction quality in both cases of the intensive and extensive margins, as indicated by the small values of the Normalized Root Mean Squared Error (NRMSE) and the Scatter Index (SI). For the intensive margin, the average difference between the predicted and observed values is 7.12 in the case of the intensive margin and 0.26 in the case of the extensive margin.

5.1.1 Products' intensive margin

Let's start by looking at the heterogeneity of the treatment effects on the intensive margin for products exported to Canada in Figure 6. We can find either products that experienced a reduction in trade following the implementation of CETA or products that consistently benefited from the new trade regime. Visually, we can realize that positive treatment effects slightly prevail. In Table 3, column (1), we report the average weighted treatment effects on the treated products, following eq. 5, which is our synthetic number to evaluate how product-level trade responded to the new trade regime. We find a positive and significant

[19]Following the original procedure by Athey *et al.* (2021), five random folds are used as cross-validation to derive the optimal λ^*_Y of eq. 7. For each λ_Y, we train our model in-sample on each k-th random training subset, $\mathcal{O}_k \subset \mathcal{O}$, and we compute $\hat{\mathbf{Y}}(\lambda_{Y(k)}, \mathcal{O}_k)$. We then calculate the RMSE for each out-of-sample k^{th} testing set. We pick the λ_Y corresponding to the minimum RMSE, which guarantees better prediction accuracy. Thus, Table 2 reports the minimum average RMSE corresponding to the optimal λ^*_Y.

value of 1.28% on export flows.[20] Interestingly, other moments of the distribution help us in evaluating the impact of the CETA. The simple average (ATET), the median, and the skewness all point to an overall positive yet asymmetric impact on product-level export flows.

Yet, the great degree of heterogeneity of the treatment effects is worth special attention, as it is a piece of evidence that has been neglected in trade policy literature. We argue that exposing heterogeneity is one important advantage of implementing matrix completion for trade policy evaluation, whereas the otherwise typical empirical test would have summarized the policy's effectiveness with a unique synthetic coefficient. For example, if we implemented a simple diff-in-diff strategy, as in Appendix B, we would obtain a unique statistically non-significant coefficient, on which we would have concluded that the treaty did not have any impact. In reality, positive and negative effects could cancel out, and the unique coefficient can conceal relevant heterogeneity.

Figure 6: Distribution of the relative Treatment Effects on the Treated (TET) - intensive margin

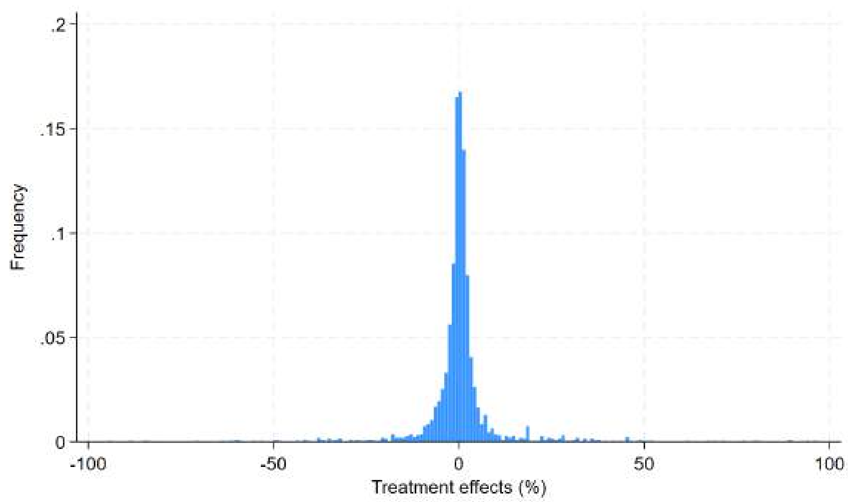

Note: The figure reports a histogram for the distribution of relative treatment effects, TET^*_{pdt}, following eq. 5, which have been computed for each HS 6-digit product exported to Canada that has seen a change in the trade regime after CETA, and then they are weighted for the relevance each product had in the year before the treaty signature.

The heterogeneity is still pronounced when we group single products by main classes, as in Table 4 and Figure 7. Apparently, most classes register a positive impact, except

[20] The statistical significance is derived from the computation of a weighted standard deviation computed as $\sqrt{\frac{\sum_{i=1}^{N} w_{pdt}\left(TET^*_{pdt} - WATET_{pdt}\right)^2}{(\mathcal{M}-1)\backslash\mathcal{M}\sum_{i=1}^{N} w_{pdt}}}$, where we take into account the distribution of weights, $\mathcal{M}$ is the number of the treatment effects on the treated products that we computed, and $WATET_{pdt}$ is the weighted average we get from 5.

Table 3: Weighted Treatment Effects on the Treated (WATET) products to Canada - intensive margin

Model	WATET (1)	weighted st. dev. (2)	N. products (3)
Intensive margin	1.278***	0.524	2,165

Note: The table reports the Weighted Average Treatment Effects on the Treated (WATET) products, obtained from TET^*_{pdt}, considering each product's relevance in the year before the treaty signature. The weighted standard deviations are computed as $\sqrt{\frac{\sum_{i=1}^{N} s_{pdt}\left(TET^*_{pdt}-WATET\right)^2}{(\mathcal{L}-1)\backslash\mathcal{L}\sum_{i=1}^{N} s_{pdt}}}$, where $\mathcal{L}$ is the number of counterfactuals in the trade matrix for Canada. Errors are bootstrapped. *** stand for $p < 0.001$.

for Animal and Animal Products, Mineral Products, Plastics/Rubbers, and Wood & Wood products; no negative impact is found on any other class. The impact is positive and higher in Foodstuffs with a weighted average treatment effect (WATET) of 1.9%, and it is lower in the case of Stone and Glass Products with a WATET of 0.476 %. Notably, almost all distributions are positively skewed with an asymmetry in favour of the positive quadrant, with the exceptions of Mineral Products (HS 25-27) and Wood & Wood Products (HS 44-49), whose WATETs are anyway non-significantly different from zero.

Nonetheless, when we evaluate the entire distribution of each product class, we always observe a fringe of products for which the signature of CETA has brought a negative impact. Even if such negative effects do not dominate the distributions, where the impact is either positive or statistically non-significant, they are still relevant and require a discussion. As a matter of fact, unweighted standard deviations are high, and they indicate huge variations around the average treatment effect. Therefore, we introduce in Section 5.1.3 a few descriptive statistics that help qualify the positive and negative variation around the albeit positive average effect.

5.1.2 Extensive margin

Figure 8 provides a snapshot of the impact on the extensive margin, while corresponding numbers are reported in Table 5. The impact is evaluated by considering the additional entry-exit dynamics due to CETA on top of the regular entry-exit that we would have seen in any case in the absence of any treatment. In Figure 8, we start by separating the exiting products on the left and the entering products on the right. The light-coloured areas indicate, in both cases, the share of entry-exit that we do not attribute to the CETA because it is regularly predicted by the matrix of potential outcomes we obtain after our algorithm. The dark-coloured area represents instead the cases of treatment effects (TET) that are different

Figure 7: Distribution of the relative Treatment Effects (TE) on the intensive margin by main product classes

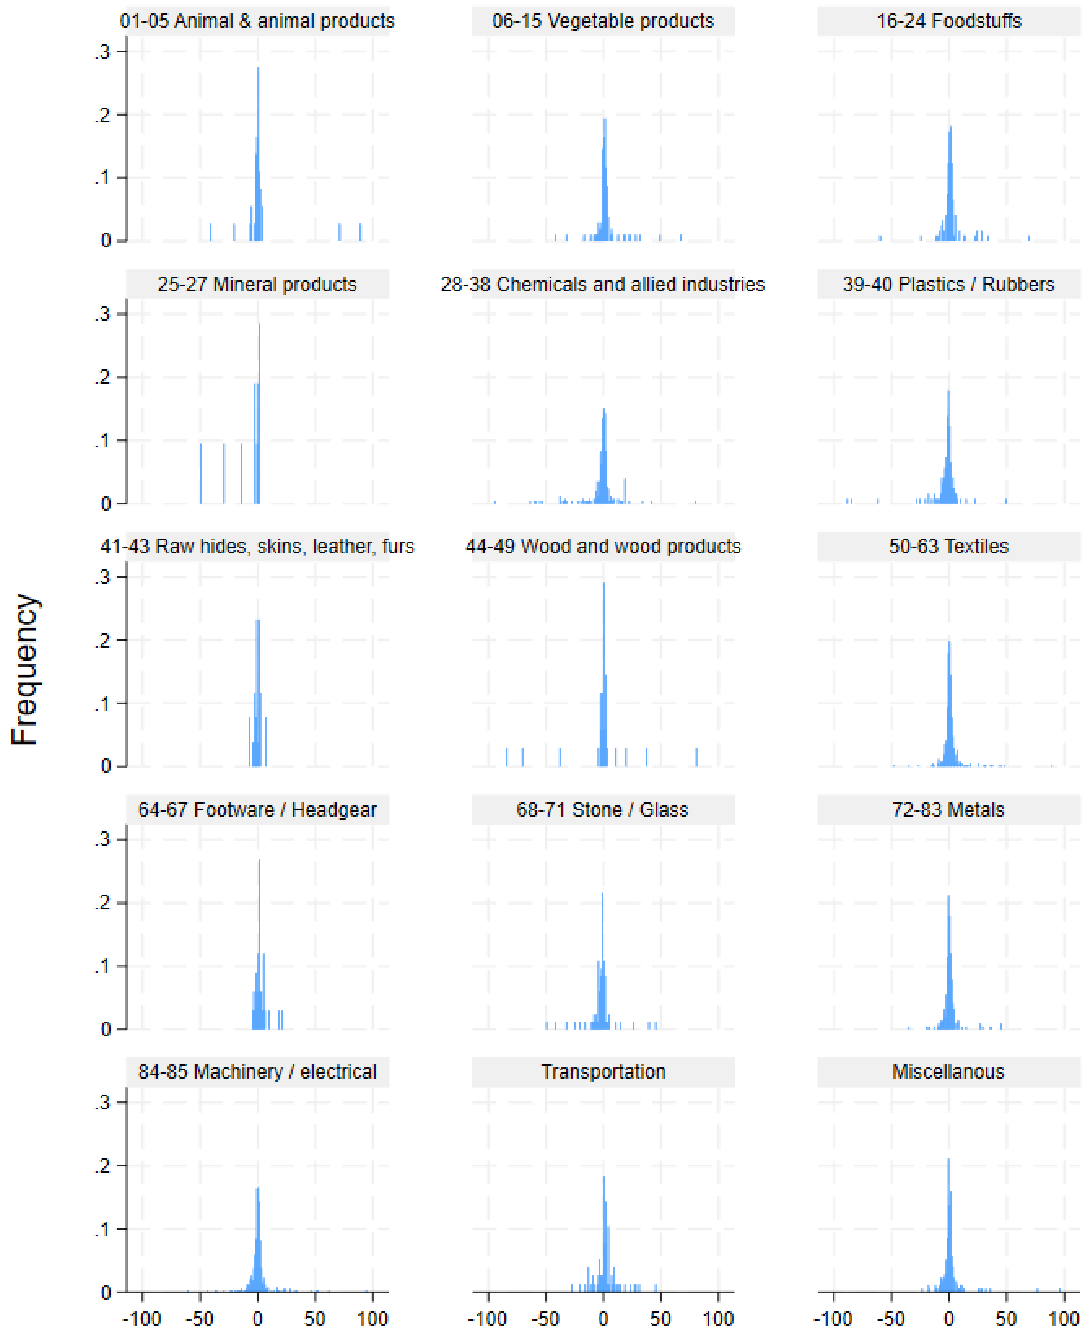

Note: The figure reports histograms for the distribution by main product classes of relative treatment effects, TE_{pdt^*}, following eq. 5, which have been computed for each HS 6-digit product exported to Canada that has seen a change in the trade regime after CETA, and then they are weighted for the relevance each product had in the year before the treaty signature.

Table 4: Weighted Average Treatment Effects on the Treated (WATET) products to Canada - intensive margin of main product classes

Product class	Class name	WATET	weighted st. dev.	N. products
01-05	Animal & Animal Products	0.503	1.341	43
06-15	Vegetable Products	0.958**	0.363	109
16-24	Foodstuffs	1.902***	0.125	130
25-27	Mineral Products	1.000	0.547	11
28-38	Chemicals & Allied Industries	1.161**	0.406	232
39-40	Plastics / Rubbers	0.454	0.498	129
41-43	Raw Hides, Skins, Leather & Furs	0.679***	0.182	27
44-49	Wood & Wood products	1.073	0.717	36
50-63	Textiles	1.351***	0.167	442
64-67	Footwear / Headgear	1.337***	0.275	36
68-71	Stone / Glass	0.476*	0.183	88
72-83	Metals	1.4*	0.620	230
84-85	Machinery / Electrical	0.927***	0.277	417
86-89	Transportation	1.249*	0.562	83
90-97	Miscellaneous	1.119***	0.239	186

Note: The table reports the Weighted Average Treatment Effects on the Treated (WATET) exports by main product classes to Canada. Treatment effects in percentage points, TET^*_{pdt}, are weighted for each product's relevance in the year before the treaty signature to obtain the unique $WATET$. The weighted standard deviations are computed as $\sqrt{\frac{\sum_{i=1}^{N} s_{pdt}\left(TET^*_{pdt}-WATET\right)^2}{(\mathcal{L}-1)\backslash\mathcal{L}\sum_{i=1}^{N} s_{pdt}}}$, where $\mathcal{L}$ is the total number of the treatment effects on the treated units for the reference population of each row. Errors are bootstrapped. *** stand for $p < 0.001$.

from zeros, as from eq. 3. If we compare with the number of incumbent products[21] in 2017; the bar on the left indicates a positive extensive margin of about 14.5%. That is, in 2017, we had an additional 14.5% of products exported from France to Canada for the first time, thanks to CETA. On the other hand, we register a negative extensive margin equal to 13.1% if we compare it with incumbent products. That is, in 2017, we had an additional 13.1% of products that were not exported anymore due to CETA.

Table 5: Positive and extensive margins - with and without CETA

	with CETA	without CETA	Total
Negative extensive margin	263	123	386
Positive extensive margin	294	106	400

Note: The table reports the numbers of exiting (first row) and entering products (second row) that we observe after the signature of the CETA. In the first column, we report the numbers of products that have entered or exited due to the CETA, i.e., they are obtained as non-zero treatment effects after the matrix of potential outcomes. In the second column, we report the numbers of products that have entered or exited not due to the CETA, i.e., they are predicted as such in the matrix of potential outcomes.

[21]We consider as incumbent the 2,031 products exported in Canada after the signature of the treaty, and that were also exported at least two years before the signature of the CETA. If we consider the demography predicted by the algorithm in the absence of the CETA, we would have about 5.2% of regular entries and 6% of regular exits. These numbers are close to what we find in entry/exit in previous years, before CETA.

Figure 8: Positive and negative extensive margin

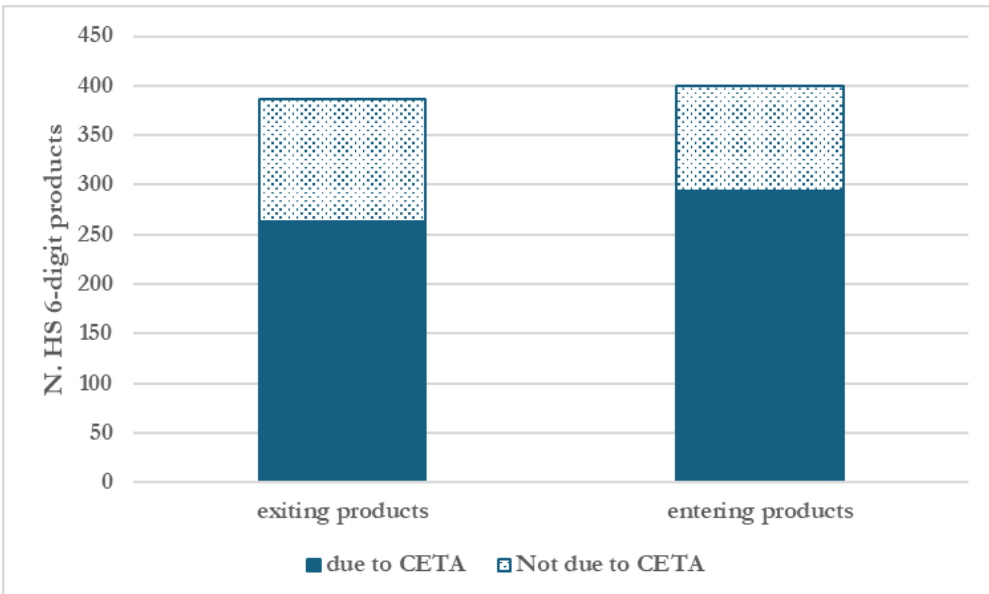

Note: The figure reports the numbers of exiting (on the left) and entering products (on the right) that we observe after the signature of the CETA. The light-coloured areas indicate products that would have entered or exited in any case without the CETA, i.e., they are predicted as such in the matrix of potential outcomes. The dark-coloured area includes products that enter or exit Canada as a result of the CETA signature, i.e., they are obtained as non-zero treatment effects after the matrix of potential outcomes.

Table 6: Extensive margin by main product classes

HS class	Product class	Exiting	Entering	Net entry
01-05	Animal & Animal Products	19	24	5
06-15	Vegetable Products	41	22	-19
16-24	Foodstuffs	6	11	5
25-27	Mineral Products	12	8	-4
28-38	Chemicals & Allied Industries	29	71	42
39-40	Plastics / Rubbers	3	1	-2
41-43	Raw Hides, Skins, Leather & Furs	1	4	3
44-49	Wood & Wood products	12	21	9
50-63	Textiles	60	31	-29
64-67	Footwear / Headgear	0	0	0
68-71	Stone / Glass	5	16	11
72-83	Metals	35	34	-1
84-85	Machinery / Electrical	31	37	6
86-89	Transportation	5	5	0
90-97	Miscellaneous	4	9	5
	Total	263	294	31

Note: The table reports the numbers of exiting (first column) and entering products (second column) by main HS product class. The focus is on the extensive margins we observe as they are due to the CETA, i.e., they are obtained as non-zero treatment effects after the matrix of potential outcomes. The third column represents the difference between the entry and the exit.

In Table 6, we further separate negative and positive extensive margins by main product classes. Here, we explicitly focus only on the entry-exit dynamics we attribute to CETA. Notably, the product class that has by far benefited the most from the treaty is the Chemicals & Allied Industries (HS 28-38), with an entry of 71 more products, followed by Machinery/Electrical (HS 84-85) with 37 new products, and Textiles (HS 50-63) with 31 new products. If we look at the negative extensive margin, we find that the group with the highest number of exits is Textiles (HS 50-63) with 60 products, followed by Vegetable Products (HS 06-15) with 41, and Metals (HS 72-83) with 35. Notably, Textiles (HS 50-63) is the class for which the net extensive margin has been most negative, with a loss of 29 products, whereas Chemicals & Allied Industries is the one with the highest gain from the net entry, with a total of 42 products.

5.1.3 Post-estimation analysis

In this section, we explore a few additional descriptive statistics that help qualify the relevant heterogeneity we detected in the previous paragraphs. We investigate the intensive and the extensive margins in Canada in relationship with a few dimensions that we deem important to describe the heterogeneity we observe.

Let us start with the results of the intensive margin. Most interestingly, we record a positive correlation between the treatment effects expressed as percentage points, TET^{*}_{pdt}, and a measure of revealed comparative advantage (RCA) computed in the year before treatment considering the universe of French customs data.[22] Eventually, in Figure 9, we visualize the statistical association with a 95% confidence interval. We observe that the correlation is positive and statistically significant after the threshold value when RCA is equal to one.

Briefly, Figure 9 shows that the higher the previous comparative advantage of the product in Canada, the higher the positive impact of the CETA. When tariffs are reduced or quotas are extended, the response in percentage points is higher for those products that were already selling well on the Canadian market. In a nutshell, a good portion of product-level heterogeneity in the effects of CETA is finally explained by initial comparative advantage positions. The latter is an interesting result that we can record because we can rely on an array of counterfactuals thanks to matrix completion.

Please note, however, that when RCA is lower than one, the association is not statistically significant. In cases of products that were at a comparative disadvantage, when a product was not selling well in Canada, it is not clear what impact we should expect after the treaty signature.

At this point, we can proceed with investigating the estimates we obtained for the extensive margin in Canada. Figure 10 reports the results of two binary regressions. In both cases, we visualize the result of a linear regression model whose dependent variable is the trade elasticity of the single HS 6-digit product sourced from Fontagné *et al.* (2022). On the left panel, a binary variable (Yes/No) declares whether the product entered the Canadian market due to the CETA or was already exported. On the right panel, a binary variable (Yes/No) declares whether the product exited the Canadian market due to the CETA or survived after the treaty. What we see is that entering and exiting products have, in general, a higher trade elasticity if compared with incumbent products. We believe it makes sense that products whose response to changes in trade costs is relatively higher are also the ones that react the most to a tariff reduction or a quota extension, eventually contributing to the extensive margin. In the case of the negative extensive margin, a fringe of exporters who face a relatively higher trade elasticity observe the changes in the relative costs and

[22] The standard measure of revealed comparative advantage (RCA) that we compute is in the form: $RCA_{pt} = \frac{\frac{X_{CA,pt}}{X_{CA,t}}}{\frac{X_{W,pt}}{X_{W,t}}}$, where $X_{CA,pt}$ is the export flow of the single p HS 6-digit product from France to Canada at time t, $X_{CA,t}$ is the total export to Canada at time t, $X_{W,pt}$ is the export of the same p product from France to the world at time t, and finally $X_{W,t}$ is the total export from France at time t.

Figure 9: Treatment Effects on the Treated (TET %) and comparative advantage - intensive margin

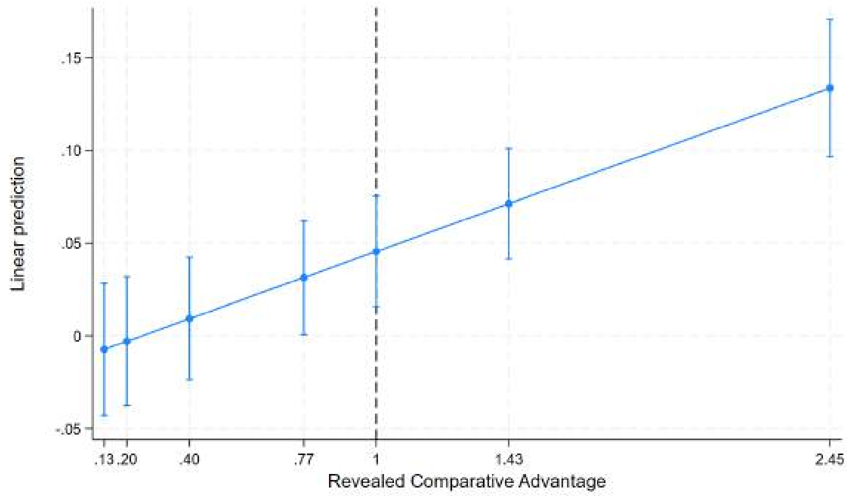

Note: The figure reports a plot of the predicted margins after a linear regression between the set of treatment effects on the treated in percentage points TET^*_{pdt} when the destination is Canada and a standard measure of Revealed Comparative Advantage computed in the year before the CETA. The reference line, when RCA is equal to one, indicates that products below it were at a comparative disadvantage and products above it were at a comparative advantage. Bars indicate a 95% confidence interval.

decide to reduce export values up to the point of exiting the Canadian market. Similarly, in the case of the positive extensive margin, a fringe of producers who face a relatively higher trade elasticity were not able to export in Canada and decided to enter the market when they observed an albeit small change in tariffs or quotas.[23]

5.2 Firm-Level Analysis

Our choice is to investigate the peculiar category of multi-product firms. The latter is an interesting category of firms that is certainly relevant, as we have seen in Figure 2 that they are responsible for about 85% of export flows from France to Canada. From another perspective, multiproduct firms are also an interesting case to follow after trade liberalization events because we can test whether they adjust their portfolios of products as predicted by trade theory.

From the original data, we select only those firms exporting more than one product to Canada within our time frame. Then, we generate a ranking for each firm by ordering products based on their trading values, from the most to the least traded by the single firm

[23] We also examined the impact of the elasticity of substitution on the intensive margin and the role of comparative advantage on the extensive margin. However, these tests did not yield any significant results.

Figure 10: Extensive margin and trade elasticity

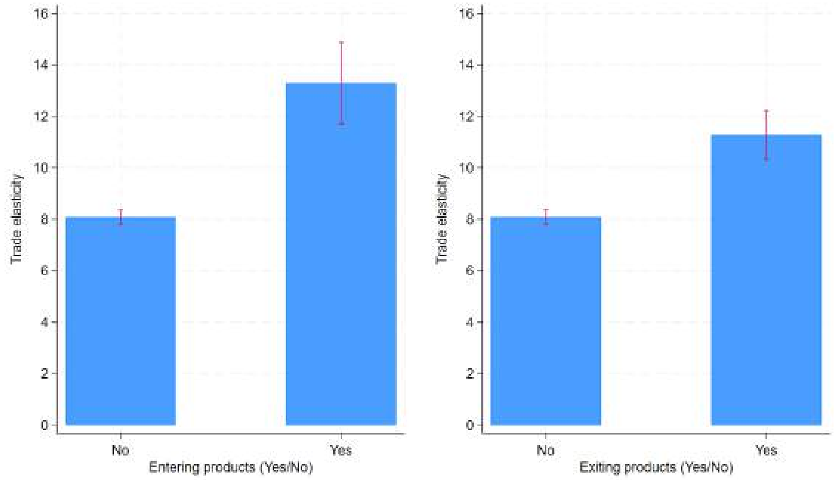

Note: The figure reports a plot of the predicted margins after two linear probability models (LPMs), whose dependent variable is the trade elasticity of the single HS 6-digit product that is exposed to the CETA. In the left panel, the comparison is between incumbent and the exiting products. In the right panel, the comparison is between the incumbent and the entering products. Trade elasticities are sourced from Fontagné *et al.* (2022). Bars indicate a 95% confidence interval.

in the year before the treaty. We will report results only on firms that trade at least three product lines to reduce the noise caused by yearly volatility in bigger portfolios of products.[24] Notably, the first most traded products at the firm level account, on average, already for 70% of that firm's exports.

In Figure 11, we report the design of a firm-level matrix to study the intensive margin by multiproduct firms. Please remember that, consistently with eq. 2, we consider as treated any (multiproduct) firm with at least one product line whose tariff or quota has been affected by the signature of the CETA. In Figure 11, rows correspond to the N multi-product French exporters. Among them, Θ is the population of treated firms, and $(N - \Theta)$ is the set of untreated firms. Each column represents a different combination of time t and product p. The product is identified at the HS 6-digit level, and we include only the three most traded lines for each firm before τ, i.e., the year of treatment, among those exported in each of the three years in the panel. The matrix element $Y_{i,(pt)}$ measures the observed outcome of firm i for the product p at time t.

Similarly to what we did at the product level, we reconstruct the matrix of observed outcomes and predict the counterfactuals following the estimator in eq. 7. Table 7 presents

[24]Please note that the exercise always needs a fixed set of products per firm to include in the matrix design. Results with two or four products per firm are available upon request. Bigger portfolios of products per firm cannot be tested for a lack of balance between the treated and the control group.

Figure 11: Matrix Structure for the firm-level analysis

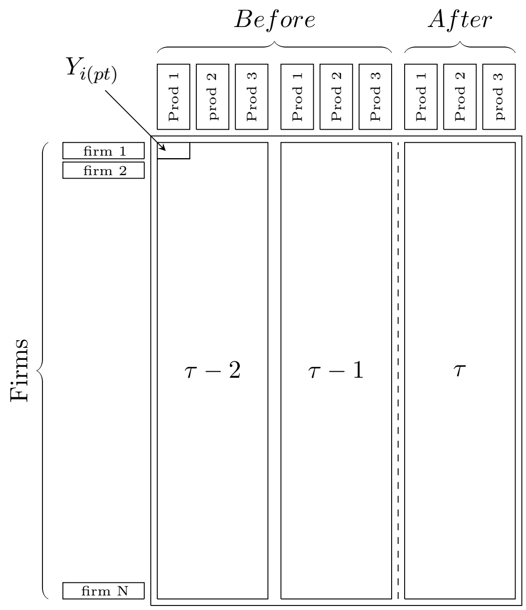

summary statistics of the prediction quality of our firm-level exercise. The percentage of expected error for the parameter of interest (i.e. the Scatter Index) is 29%. Prediction power indicates that the algorithm successfully replicates the dynamics of the original matrices of outcomes in the observed entries.[25] At this point, we can validly use predicted values of unobserved potential outcomes as counterfactuals for what would have happened if CETA was not signed.

Table 7: Prediction quality - Firm-level analysis

Model	n. obs.	$\overline{Y}$	min Av(RMSE)	SI	NRMSE
Intensive	3,177	203,345.61	59,069.2	29.04	42.93

Note: The table collects quality indicators for the predictions of observed values in the multiproduct firm-level exercise. The following columns indicate the average predicted value, the root mean squared error (RMSE), the scatter index, and the normalized RMSE.

5.2.1 Multiproduct firms and product scope

Results on the impact of CETA on multiproduct firms are reported in Table 8, while Figure 12 reports a visualization of the distributions of treatment effects for the first, second and third

[25] As in a classic machine-learning predictive framework, the algorithm is first trained on different in-sample subsets and then tested out of the sample. See also footnote 19 for further details.

exported products, respectively. Please note that, in these paragraphs, we are considering the multiproduct firms exposed to CETA and that exported at least three products in Canada vs. a control group of untreated firms, as described in eq. 4.1. Therefore, our quantities of interest are the treatment effects on the treated, TET^*_{ipt}, expressed in percentage points with reference to products ordered, $p = \{1, 2, 3\}$, after considering their trade values in the firm's portfolio before CETA.

Table 8: Weighted Average Treatment Effects on the Treated (WATET) products ranked by the multiproduct firms

Type of firm/product	WATET	weighted st. dev	N. obs
All firms			
First exported product	0.886*	0.481	418
Second exported product	0.001	0.001	418
Third exported product	0.012***	0.001	418
Manufacturing firms			
First exported product	0.729***	0.296	298
Second exported product	-0.025***	0.001	298
Third exported product	0.001	0.001	298
Trade intermediaries			
First exported product	0.157***	0.003	120
Second exported product	0.027***	0.001	120
Third exported product	0.011***	0.001	120

Note: The table reports the Weighted Average Treatment Effects on the Treated (WATET) exports for the first, second and third products in the multiproduct firms' portfolio. The $WATET$'s are computed considering products' trade shares in the year before the CETA. The weighted standard deviations are computed as $\sqrt{\frac{\sum_{i=1}^{N} s_{ipt}\left(TET^*_{ipt} - WATET\right)^2}{(\mathcal{L}-1)\backslash\mathcal{L}\sum_{i=1}^{N} s_{ipt}}}$, where $\mathcal{L}$ is the total number of the treatment effects on the treated units for the reference population of each row. *, **, *** stand, respectively, for $p < 0.05$, $p < 0.01$, $p < 0.001$.

If we look at the first part of Table 8, we find that the weighted average treatment effect on the treated (WATET) first products is 0.87%, although weakly significant. At the same time, the WATET on the second product is not significantly different from zero, while the WATET on the third product indicates a tiny yet significant increase of 0.012%. Briefly, the CETA has, on average, a positive impact on at least two products out of three in

Figure 12: Distribution of treatment effects (%) by product ranked in multiproduct firms

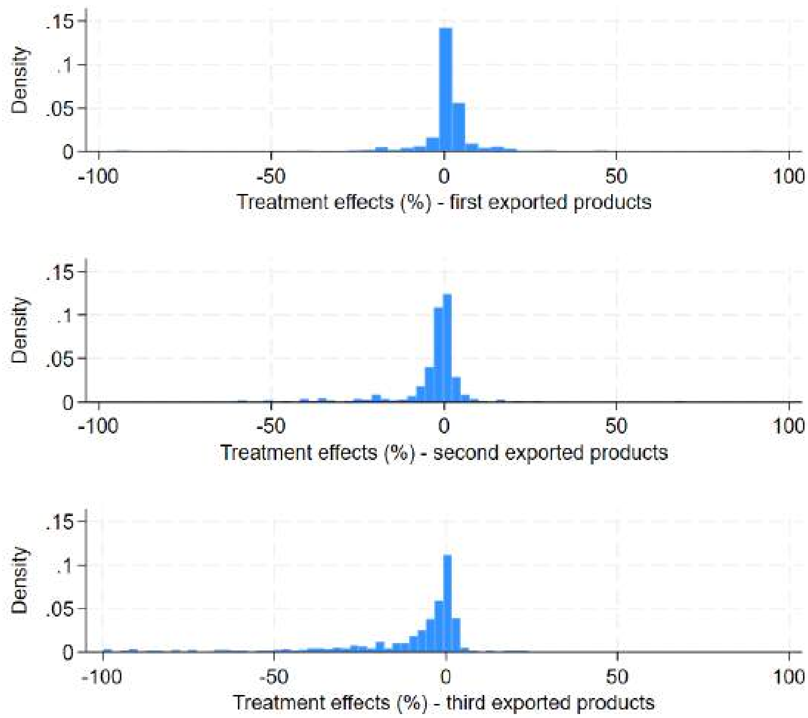

Note: The Table shows the distribution of the treatment effects on the treated in percentage points, TET^{*}_{ipt}, for the first, second and third exported products in the multiproduct firms' portfolio.

the portfolio of multiproduct firms exposed to CETA. Yet, the impact is bigger for products already performing better in the Canadian market. Visually, our results are confirmed by the three graphs we included in Figure 12 where, however, we can observe relevant heterogeneity in the positive and negative quadrants.

Importantly, the second and third parts of Table 8 differentiate firms separating manufacturing firms from those firms that professionally act as intermediaries on behalf of other firms.[26] Our separation is based on the NACE rev. 2 core activities of the firms, on which we assume that wholesalers and retailers (NACE 45, 46 and 47) work as trade intermediaries in our data. It is interesting to see that, in the case of manufacturing firms, the first exported products sell about 0.73% more, whereas the second exported products sell an almost negligible 0.03% less after the CETA. When we look at trade intermediaries, we confirm that the impact on exported products is, on average, higher, but we still find positive albeit minor effects on second and third products.

[26] Originally, our data also included firms in primary markets, like agricultural products and other commodities, in the NACE rev. 2 sectors 01-09. However, none of these firms are multiproduct if we follow the definition we introduced, and they are excluded from this part of the analysis.

Finally, we believe previous results are in line with a mechanism of portfolio adjustment predicted by trade theory, as in Mayer *et al.* (2014) and Eckel & Neary (2010). According to trade theory, liberalization events also entail more competition in an export market. More firms can access the Canadian market, and competitive pressure induces exporters to concentrate their efforts on their best-performing products, thus focusing on their core competencies. Our findings are also confirmed by a quick check on aggregate flows. According to our data, after trade liberalization between Canada and France with CETA, the first products by French exporters concentrated about 77% of the total firms' exports, which is an increase with respect to a share of 70% registered just before the treaty signature.

5.3 General equilibrium trade impacts

Our approach allows us to consider destinations different from Canada and, hence, to test whether CETA has brought about any trade diversion effects. In fact, the product-level matrix we designed in Figure 5 included alternative destinations, of which ten top partners of France and the rest are continental aggregates[27], while we always have considered the treatment to be product-specific, to have the possibility to evaluate what happens in the destinations alternative to Canada. As a consequence, our matrix completion algorithm returns us counterfactuals on sixteen destinations, including Canada, and we can check the treatment effects on the treated, TET^{*}_{pdt}, for each HS 6-digit product p exposed to the CETA, which is exported to a destination d different from Canada in time t.

The mechanism is that any trade liberalization event, including CETA, changes the distribution of relative costs incurred by exporters. A tariff decrease in Canada increases the relative cost of exporting to other destinations. This is especially true when we are in the presence of bigger exporters, who can adjust their portfolio of destinations once they internalize the new distribution of relative costs across the globe. Eventually, this is the classical Vinerian diversion effect Viner (1950), whereby trade between parties to a PTA partially substitutes for trade between partied and third countries.

We test this mechanism by estimating the following model:

$$TET_{dpt} = \alpha + \beta TET_{CA,pt} + \gamma Export\ Value_{dp,t-1} + \eta_{dpt} \tag{8}$$

where the dependent variable is the treatment effects on the treated products expressed in monetary values, TET_{dpt}, with d different from Canada; $TET_{CA,pt}$ is the treatment effects on the same treated products in Canada and $Export\ Value_{dp,t-1}$ is the initial value of the

[27]The complete list is reported in Table A3. Alternative trade destinations have been picked considering a combination of two ranks: export values and numbers of exported products.

trade flows in the alternative destination $d \neq CA$. We report the estimated results in Table 9. This time, we consider treatment effects in monetary values because we want to check whether there is a correlation in the magnitudes with the treatment effects on the treated in Canada, $TET_{CA,pt}$. Our coefficients of interest are on the first row. When we control for the initial value of the trade flows in the alternative destination (column 2), we find a negative association equal to 1.042 between the export change in Canada and the export changes of the same products in the alternative destinations. This association is robust to the inclusion of a double clustering of errors by country and by product classes (column 3). Notably, when we separate between products by their trade elasticity sourced from Fontagné *et al.* (2022), we discover that the association is mainly driven by the most elastic products (column 5), i.e., the ones whose elasticity value is above the median computed on the entire distribution. Briefly, export flows of products listed by the CETA adjust in alternative destinations as a consequence of the expected general equilibrium effects. We believe the latter is a powerful result that confirms the existence of mechanisms of reallocation on a global scale, as in the case of trade diversion, to take into account the changing distribution of relative trade costs after a liberalization event.

Table 9: CETA and alternative destinations - general equilibrium trade effects

Dependent variable TET_{pdt}	(1)	(2)	(3)	(4)	(5)
$TET_{CA,pt}$	-0.552	-1.042**	-1.042***	-0.101	-1.745***
	(0.449)	(0.437)	(0.364)	(0.154)	(0.639)
$Value_{pdt-1}$		0.019***	0.019***	0.006***	0.019***
		(0.004)	(0.002)	(0.002)	(0.001)
Constant	52,182.54***	-56,199.12**	-56,199.12***	-10,115.56*	-51,568.57***
	(11,497.63)	(24,598.21)	(17,027.99)	(4,845.29)	(16,505)
N. obs.	31,758	31,758	31,758	15,445	16313
R squared	0.0012	0.8123	0.8123	0.1890	0.8294
Clusters by country	No	Yes	Yes	Yes	Yes
Clusters by product class	No	No	Yes	Yes	Yes
Elasticity of subst.	All	All	All	below median	above median

Note: The Table shows results after a linear regression model whose dependent variable includes the treatment effects on the treated in monetary values, TET_{pdt}, where destination d is different from Canada. The main regressor of interest is the vector of treatment effects on the treated in monetary values, TET_{pdt}, where destination d is instead Canada. The unique control variable is the value of the product p export flow in destination d different from Canada in the period before the CETA, $t-1$. Errors are double-clustered by country and product class. Trade elasticity is sourced from Fontagné *et al.* (2022) **, *** stand, respectively, for $p < 0.05$, $p < 0.01$, $p < 0.001$.

Finally, we want to check whether trade diversion effects are heterogeneous by country. Thus, we reproduce the exercise from eq. 8 for separate alternative destinations, and we report results for the top ten in Table 10. Most interestingly, the trade diversion effect is higher for exports in the United States. The aggregate of African countries comes second; then, we find Germany and Italy. Exports to China, Spain, Belgium, and the United Kingdom are less responsive, while trade diversion from Netherlands and Switzerland is not statistically significant.

Table 10: Top 10 destinations - general equilibrium trade effects

Destinations	Coef.	std. err.	R squared	N. obs.
United States	-1.315**	(.565)	0.943	2,099
Africa	-1.163**	(.438)	0.372	2,153
Germany	-1.137**	(.429)	0.893	2,146
Italy	-1.006**	(.385)	0.481	2,134
China	-.647***	(.151)	0.993	1,996
Spain	-.636**	(.235)	0.629	2,143
Belgium	-.503**	(.204)	0.443	2,146
United Kingdom	-.460*	(.217)	0.608	2,136
Netherlands	-0.301	(0.234)	0.571	2,113
Switzerland	-.0781	(.179)	0.314	2,146

Note: The Table shows results after a linear regression model whose dependent variable includes the treatment effects on the treated in monetary values, TET_{pdt}, where d is one of the top 10 destinations of French exports. The main regressor of interest is the vector of treatment effects on the treated in monetary values, TET_{pdt}, where destination d is instead Canada. The unique control variable is the value of the product p export flow in destination d different from Canada in the period before the CETA, $t-1$. Errors are clustered by product class. **, *** stand, respectively, for $p < 0.05$, $p < 0.01$, $p < 0.001$.

6 Robustness and sensitivity checks

Our first concern is that products could have been endogenously selected by the parties during the treaty negotiations, and we may pick a positive impact just because selected products already showed a higher trade potential. Clues of an endogenous selection into the treaty were offered in Table 1. Products in the CETA were already exported by a greater number of French firms, more frequently, with a lower average transaction value and a lower average value dispersion. To address this concern, we conduct a placebo test by replicating the matrix completion analysis using the same definition of treated products as in the baseline, but for the period September 2012-August 2015. In Appendix Table A4, we report no significant effect, and we argue that this is supporting evidence for our empirical approach, which is capable of handling cherry-picking selections into the treaty.

A second concern is that specific matrix configurations can drive different results. The concern is specifically relevant to the validity of our findings on trade diversion when we search for possible general equilibrium effects. In this case, we test different configurations for how destinations alternative to Canada are included in the baseline matrix. In Appendix Table A5, we show results when:

1. we consider the popularity of alternative destinations classified by the number of French exporters that serve them;

2. we adopt a measure of import structure similarity to Canada, computed considering the sums of the absolute values of the distances between the share of each product p in destination d and the corresponding share of imports in Canada;

3. we select destinations based on the size of their import market.

Interestingly, the baseline estimates of the WATET for the products' intensive margin consistently fall in an interval $[0.94, 1.22]$, which is only slightly lower than our baseline estimates at 1.28%. Importantly, Appendix Table A6 confirms also the robustness of general equilibrium effects when we select destinations based on either the number of French exporters or the size of the import market. When we consider similar import structures to Canada, the coefficient of interest is not statistically significant anymore, and we argue that it makes sense because the selected destinations are less relevant for French exporters. Notably, none of the alternative matrix configurations[28] achieved the same level of prediction performance as our baseline, as shown in Appendix Table A7. For this reason, we prefer to keep our baseline matrices.

A third concern is that results are driven by the specific choice of a matrix completion algorithm. As we discussed in Section 4, the main difference between the algorithm that we adapt from Athey *et al.* (2021) and standard proposals in computer science literature (Candes & Plan, 2010; Candes & Recht, 2012) is the inclusion of vectors of fixed effects before proceeding with the singular value decomposition. In our case, we remove the vector of firm-level fixed effects, and we find that the prediction performance slightly worsens. We do not see a fundamental change in the results, but we prefer to keep our baseline results.

Finally, we investigate what happens when we change the definition of treated firms. In our baseline, a multiproduct firm is treated when it exports at least two products in Canada and, among them, at least one is enlisted by the CETA. Briefly, by our definition, we have some treated firms with a portfolio that encompasses both products that have seen a regime change and products that have not. If we change our definition and consider as treated only those firms that export at least two products all enlisted by the CETA, what we observe is that the sample shrinks dramatically to the point that it is not representative anymore. In fact, we have that 41% of multiproduct firms usually have in their portfolio both product types; they are usually bigger exporters, and we would introduce a major sample selection. For this reason, we conclude that results with a different definition of treated firms cannot be trusted.

[28]The list of alternative destinations by each selection strategy is reported in Appendix Table A8.

7 Conclusions

The present work proposes a novel approach to evaluating the impact of trade agreements using a causal machine learning framework. The aim is to provide a robust empirical strategy capable of handling the complexities and heterogeneity of trade effects at both the product and firm levels while mitigating concerns about endogenous selections into trade agreements. As a case study, we consider the entry into force of the EU-Canada Comprehensive Economic and Trade Agreement (CETA) and adapt an algorithm proposed by Athey *et al.* (2021) to the case of French customs data. The main advantage is that we can predict multidimensional counterfactuals at the firm, product, and destination levels and, thus, obtain consistent estimates of causal effects.

Findings reveal an average small albeit statistically significant positive impact of the CETA on the product-level intensive margin in the year after the CETA. The Weighted Average Treatment Effects on the Treated (WATET) is 1.28%. Yet, product-level heterogeneity of the impact is relevant, and we show how the full distribution of treatment effects needs to be evaluated. Notably, we find that the impact is higher on those products for which France showed a comparative advantage before the treaty. On the extensive margin, we record a product churning due to the treaty, which goes beyond the numbers of regular entry-exit dynamics. Due to the CETA agreement, there is a 13.1% of products not exported before that substitute 11.9% of products that are no longer exported. Interestingly, entering and exiting products are those that are more responsive to trade cost changes, i.e., whose trade elasticity is higher. At the firm level, we test the case of multiproduct firms. Consistent with the mechanism of portfolio adjustment predicted by Mayer *et al.* (2014), we observe that multiproduct exporters reallocate shares towards their first and most exported product, possibly due to an increasing local market competition after trade liberalization. Finally, our empirical strategy is suitable for capturing general equilibrium effects. Indeed, when we look at alternative destinations, we show that CETA induces trade diversion. As the trade treaty makes destinations different from Canada relatively more costly, product flows are partly redirected from other destinations towards Canada.

In conclusion, we believe we showed the validity of a matrix completion approach in evaluating changing trade policies. We believe that while the specific results have limited external validity, as they depend on the specific nature of French Trade, the same approach can be adapted in the evaluation of other trade policy actions. The main advantage is the possibility of predicting multidimensional counterfactuals as cells of a well-designed matrix, thus returning a more complete picture of the heterogeneity of the impact of trade regime changes, including general equilibrium effects from different policies in destinations that are not parties to trade agreements.

References

ABADIE, ALBERTO, DIAMOND, ALEXIS, & HAINMUELLER, JENS. 2010. Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American statistical Association*, **105**(490), 493–505.

ABADIE, ALBERTO, DIAMOND, ALEXIS, & HAINMUELLER, JENS. 2015. Comparative Politics and the Synthetic Control Method. *American Journal of Political Science*, **59**(2), 495–510.

ABRAMS, RICHARD K. 1980. International trade flows under flexible exchange rates. *Economic Review*, **65**(3), 3–10.

AITKEN, NORMAN D. 1973. The effect of the EEC and EFTA on European trade: A temporal cross-section analysis. *The American Economic Review*, **63**(5), 881–892.

ANDERSON, JAMES E, & YOTOV, YOTO V. 2016. Terms of trade and global efficiency effects of free trade agreements, 1990–2002. *Journal of International Economics*, **99**, 279–298.

ANGRIST, JOSHUA D., & PISCHKE, JÖRN-STEFFEN. 2010. The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics. *Journal of Economic Perspectives*, **24**(2), 3–30.

ARKHANGELSKY, DMITRY, ATHEY, SUSAN, HIRSHBERG, DAVID A, IMBENS, GUIDO W, & WAGER, STEFAN. 2019. *Synthetic Difference in Differences*. Tech. rept. National Bureau of Economic Research.

ARKOLAKIS, COSTAS, & MUENDLER, MARC-ANDREAS. 2013. Exporters and their products: a collection of empirical regularities. *CESifo Economic Studies*, **59**(2), 223–248.

ARKOLAKIS, COSTAS, GANAPATI, SHARAT, & MUENDLER, MARC-ANDREAS. 2021. The extensive margin of exporting products: A firm-level analysis. *American Economic Journal: Macroeconomics*, **13**(4), 182–245.

ATHEY, SUSAN, & IMBENS, GUIDO W. 2019. Machine Learning Methods That Economists Should Know About. *Annual Review of Economics*, **11**(1), 685–725.

ATHEY, SUSAN, BAYATI, MOHSEN, DOUDCHENKO, NIKOLAY, IMBENS, GUIDO, & KHOSRAVI, KHASHAYAR. 2021. Matrix completion methods for causal panel data models. *Journal of the American Statistical Association*, 1–15.

BAIER, SCOTT L, & BERGSTRAND, JEFFREY H. 2002. *On the endogeneity of international trade flows and free trade agreements*. Tech. rept. mimeo New York.

Baier, Scott L., & Bergstrand, Jeffrey H. 2004. Economic determinants of free trade agreements. *Journal of International Economics*, **64**(1), 29–63.

Baier, Scott L, & Bergstrand, Jeffrey H. 2007. Do free trade agreements actually increase members' international trade? *Journal of international Economics*, **71**(1), 72–95.

Baier, Scott L, & Bergstrand, Jeffrey H. 2009. Estimating the effects of free trade agreements on international trade flows using matching econometrics. *Journal of international Economics*, **77**(1), 63–76.

Baier, Scott L., Yotov, Yoto V., & Zylkin, Thomas. 2019. On the widely differing effects of free trade agreements: Lessons from twenty years of trade integration. *Journal of International Economics*, **116**, 206–226.

Baldwin, John, & Gu, Wulong. 2009. The impact of trade on plant scale, production-run length and diversification. *Pages 557–592 of: Producer dynamics: New evidence from micro data.* University of Chicago Press.

Bas, Maria, & Bombarda, Pamela. 2013. Chinese trade reforms, market access and foreign competition: The patterns of french exporters. *the world bank economic review*, **27**(1), 80–108.

Bergstrand, Jeffrey H. 1985. The gravity equation in international trade: some microeconomic foundations and empirical evidence. *The review of economics and statistics*, 474–481.

Bernard, Andrew B, Redding, Stephen J, & Schott, Peter K. 2010. Multiple-product firms and product switching. *American economic review*, **100**(1), 70–97.

Bernard, Andrew B, Redding, Stephen J, & Schott, Peter K. 2011. Multiproduct firms and trade liberalization. *The Quarterly journal of economics*, **126**(3), 1271–1318.

Candes, Emmanuel, & Recht, Benjamin. 2012. Exact matrix completion via convex optimization. *Communications of the ACM*, **55**(6), 111–119.

Candes, Emmanuel J, & Plan, Yaniv. 2010. Matrix completion with noise. *Proceedings of the IEEE*, **98**(6), 925–936.

Chernozhukov, Victor, Wüthrich, Kaspar, & Zhu, Yinchu. 2021. An Exact and Robust Conformal Inference Method for Counterfactual and Synthetic Controls. *Journal of the American Statistical Association*, **116**(536), 1849–1864.

Crowley, Meredith, & Han, Lu. 2022 (7). *The Pro-competitive Effects of Trade Agreements.* CEPR Discussion Paper No. 17463.

De Chaisemartin, Clément, & d'Haultfoeuille, Xavier. 2023. Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey. *The Econometrics Journal*, **26**(3), C1–C30.

Dhingra, Swati. 2013. Trading away wide brands for cheap brands. *American Economic*

Review, **103**(6), 2554–84.

Eaton, Jonathan, & Fieler, Ana Cecília. 2019. *The margins of trade*. Tech. rept. National Bureau of Economic Research.

Eaton, Jonathan, Kortum, Samuel, & Kramarz, Francis. 2004. Dissecting trade: Firms, industries, and export destinations. *American Economic Review*, **94**(2), 150–154.

Eaton, Jonathan, Kortum, Samuel, & Kramarz, Francis. 2011. An anatomy of international trade: Evidence from French firms. *Econometrica*, **79**(5), 1453–1498.

Eaton, Jonathan, Kortum, Samuel S, & Sotelo, Sebastian. 2012. *International trade: Linking micro and macro*. Tech. rept. National bureau of economic research.

Eckel, Carsten, & Neary, J Peter. 2010. Multi-product firms and flexible manufacturing in the global economy. *The Review of Economic Studies*, **77**(1), 188–217.

Egger, Peter H., & Tarlea, Filip. 2021. Comparing Apples to Apples: Estimating Consistent Partial Effects of Preferential Economic Integration Agreements. *Economica*, **88**(350), 456–473.

Feenstra, Robert, & Ma, Hong. 2007. *Optimal choice of product scope for multiproduct firms under monopolistic competition*. Tech. rept. National Bureau of Economic Research.

Feenstra, Robert C, Markusen, James R, & Rose, Andrew K. 2001. Using the gravity equation to differentiate among alternative theories of trade. *Canadian Journal of Economics/Revue canadienne d'économique*, **34**(2), 430–447.

Fontagné, Lionel, Secchi, Angelo, & Tomasi, Chiara. 2018. Exporters' product vectors across markets. *European Economic Review*, **110**, 150–180.

Fontagné, Lionel, Guimbard, Houssein, & Orefice, Gianluca. 2022. Tariff-based product-level trade elasticities. *Journal of International Economics*, **137**, 103593.

Goldberg, Pinelopi K, & Pavcnik, Nina. 2016. The effects of trade policy. *Pages 161–206 of: Handbook of commercial policy*, vol. 1. Elsevier.

Head, Keith, & Mayer, Thierry. 2014. Gravity equations: Workhorse, toolkit, and cookbook. *Pages 131–195 of: Handbook of international economics*, vol. 4. Elsevier.

Head, Keith, & Ries, John. 1998. Immigration and trade creation: econometric evidence from Canada. *Canadian journal of economics*, 47–62.

Hübner, Kurt, Deman, Anne-Sophie, & Balik, Tugce. 2017. EU and trade policy-making: the contentious case of CETA. *Journal of European Integration*, **39**(7), 843–857.

Iacovone, Leonardo, & Javorcik, Beata S. 2010. Multi-product exporters: Product churning, uncertainty and export discoveries. *The Economic Journal*, **120**(544), 481–499.

Larch, Mario, & Yotov, Yoto V. 2023 (2). *Estimating the Effects of Trade Agreements: Lessons from 60 Years of Methods and Data*. CESifo Working Paper No. 10267.

Limão, Nuno. 2016. Preferential trade agreements. *Pages 279–367 of: Handbook of com-*

mercial policy, vol. 1. Elsevier.

MAGEE, CHRISTOPHER S. 2003. Endogenous preferential trade agreements: An empirical analysis. *Contributions in Economic Analysis & Policy*, **2**(1), 1–17.

MAYER, THIERRY, MELITZ, MARC J, & OTTAVIANO, GIANMARCO IP. 2014. Market size, competition, and the product mix of exporters. *American Economic Review*, **104**(2), 495–536.

MAYER, THIERRY, MELITZ, MARC J, & OTTAVIANO, GIANMARCO IP. 2021. Product mix and firm productivity responses to trade competition. *Review of Economics and Statistics*, **103**(5), 874–891.

MAZUMDER, RAHUL, HASTIE, TREVOR, & TIBSHIRANI, ROBERT. 2010. Spectral regularization algorithms for learning large incomplete matrices. *The Journal of Machine Learning Research*, **11**, 2287–2322.

MELITZ, MARC J. 2003. The impact of trade on intra-industry reallocations and aggregate industry productivity. *econometrica*, **71**(6), 1695–1725.

MELITZ, MARC J, & OTTAVIANO, GIANMARCO IP. 2008. Market size, trade, and productivity. *The review of economic studies*, **75**(1), 295–316.

MRÁZOVÁ, MONIKA, & NEARY, J PETER. 2017. Not so demanding: Demand structure and firm behavior. *American Economic Review*, **107**(12), 3835–3874.

MULLAINATHAN, SENDHIL, & SPIESS, JANN. 2017. Machine Learning: An Applied Econometric Approach. *Journal of Economic Perspectives*, **31**(2), 87–106.

NAGENGAST, ARNE, & YOTOV, YOTO V. 2024. Staggered difference-in-differences in gravity settings: Revisiting the effects of trade agreements. *American Economic Journal: Applied Economics*.

QIU, LARRY D, & ZHOU, WEN. 2013. Multiproduct firms and scope adjustment in globalization. *Journal of International Economics*, **91**(1), 142–153.

RUBIN, DONALD B. 2005. Causal Inference Using Potential Outcomes. *Journal of the American Statistical Association*, **100**(469), 322–331.

SOLOAGA, ISIDRO, & WINTERSB, L ALAN. 2001. Regionalism in the nineties: What effect on trade? *The North American Journal of Economics and Finance*, **12**(1), 1–29.

VINER, JACOB. 1950. *The customs union issue*. Carnegie Endowment for International Peace, New York.

WESTERLUND, JOAKIM, & WILHELMSSON, FREDRIK. 2011. Estimating the gravity model without gravity using panel data. *Applied Economics*, **43**(6), 641–649.

YANG, SHANPING, & MARTINEZ-ZARZOSO, INMACULADA. 2014. A panel data analysis of trade creation and trade diversion effects: The case of ASEAN–China Free Trade Area. *China Economic Review*, **29**, 138–151.

Appendix A: Tables and graphs

Table A1: Distribution of tariff changes in the Canada-EU Comprehensive Trade Agreement (CETA)

Tariff decrease (%)	N. products	% products
0.3 - 5	1,871	51.04
6 - 10	1,290	35.19
11 - 20	479	13.06
>20	26	0.71
Total	3,666	100.00

Note: The table shows the distribution of tariff changes by HS 6-digit products as it has been negotiated in the CETA. The simple average tariff decrease has been 5.8% with a 4.3 standard deviation.

Figure A1: Products per exporter in Canada in 2016

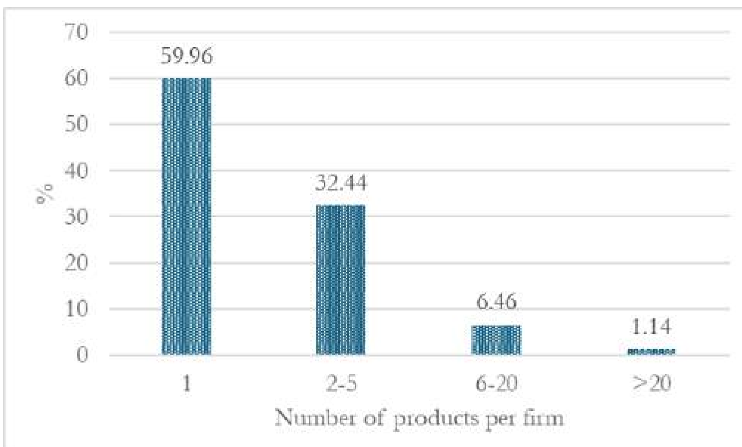

Note: The figure shows the distribution of product portfolios by exporters to Canada before the entry into force of the CETA. On the left, the first bar indicates exporters with one product delivered to Canada. Then, the following bars refer to product portfolios sold to Canada by multiproduct firms.

Table A2: Which products in the intensive margin

Case	Traded in 2015	Traded in 2016	Traded in 2017	Intensive margin	Note:
1)	Yes	Yes	Yes	Yes	Always traded
2)	Yes	Yes	No	No	Not traded after CETA
3)	Yes	No	Yes	Yes	Intermittently traded
4)	No	Yes	Yes	Yes	Intermittently traded
5)	Yes	No	No	No	Intermittently traded
6)	No	Yes	No	Yes	Intermittently traded
7)	No	No	Yes	No	Traded only after CETA
8)	No	No	No	No	Never traded

Note: The table separates cases of intensive margins from different trade patterns in the original data. For each of them, we report in column (4) whether the corresponding product is included in the analyses on intensive margins.

Table A3: Ranking export destination by trade volumes and number of products

Destination	Export volume (in mln €) (1)	# Products (2)	Rank by Values (3)	Rank by #Products (4)	Combined Rank (5)
Germany	73.134	4,816	1	4	2.5
Italy	36.084	4,842	4	2	3
Spain	36.692	4,825	3	3	3
Belgium	30.752	4,857	6	1	3.5
USA	38.771	4,091	2	9	5.5
United Kingdom	35.721	4,594	5	7	6
Netherlands	16.350	4,775	8	5	6.5
Switzerland	15.922	4,691	9	6	7.5
China	19.489	3,836	7	10	8.5
Poland	8.356	4,193	10	8	9
Canada	4.217	3,812	26	18	22
Rest of Asia	74.958	4,890			
Rest of Europe	55.077	4,999			
Africa	29.825	4,912			
Rest of Americas	16.377	4,245			
Oceania	5.449	4,106			

Note: Countries in this table are included in the trade matrix at the product level introduced in Section 5.1. The decision is based on two criteria: in column (1), we report the average trade values exported by the French Exporters from 2015-2016; in column (2), we report the average number of products exported to each destination in 2015-2016. Columns (3) and (4) report the ranking position of each country by average trade values and average number of exported products, respectively. Column (5) reports an average of rankings in columns (3) and (4). The (rest of the) continents at the bottom of the table are also included in the analyses to close and balance the trade matrix.

Figure A2: Time-destination correlation matrices on the extensive margins: original data matrix and the predicted low-rank matrix

(a) Original data matrix $\mathbf{Y}$ (b) Predicted matrix $\tilde{\mathbf{Y}}$

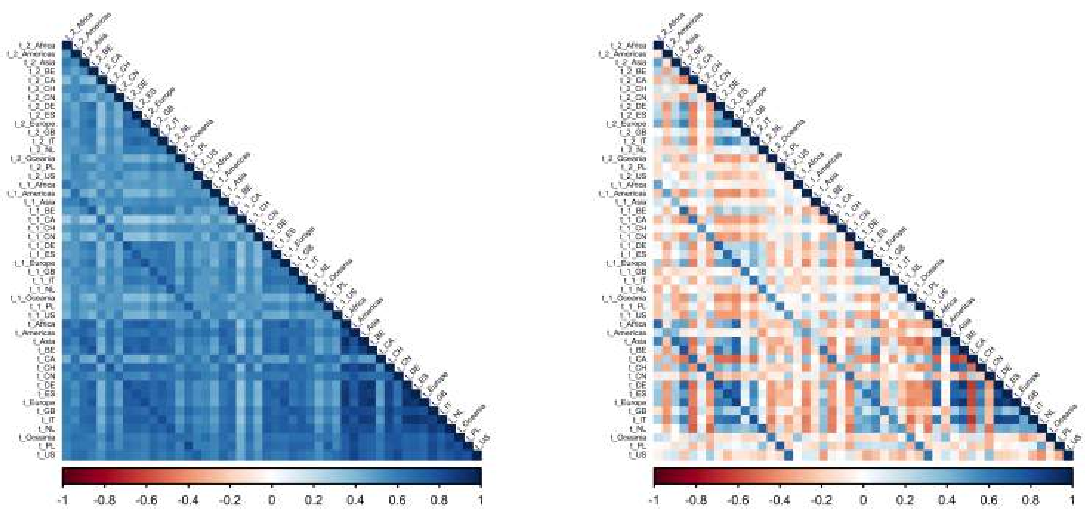

Note: he figure on the left displays the correlation matrix of the columns of the original matrix $\mathbf{Y}$, which records the exports of products (at the HS6 level) to country d at time t. The figure on the right shows the correlation matrix of the columns of the corresponding predicted low-rank matrix $\tilde{\mathbf{Y}}$. This predicted matrix accounts for the residual correlation between the rows and columns of the original matrix after removing the row and column fixed effects ($\tilde{\gamma}_i$ and $\tilde{\delta}_j$ in equation 7). Figure (a) illustrates significant and consistent correlation patterns between destinations over time. Figure (b) demonstrates that these correlation patterns are effectively learned and captured by the predicted low-rank matrix $\tilde{\mathbf{Y}}$, enhancing the accuracy of the predicted outcomes.

Table A4: A placebo test for the intensive margin to Canada

Product class	Class name	WATET (1)	weighted st. dev. (2)	N. products (3)
01-97	All products	-1.038	11.664	2,219
01-05	Live animals & Animal products	0.932	85.550	44
06-15	Vegetable products	5.380	0.696	122
16-24	Foodstuffs	0.415	4.262	120
25-27	Mineral products	-32.675	232.346	23
28-38	Chemicals & Allied industries	-1.613	12.084	244
39-40	Plastics / Rubbers	-1.289	9.967	129
41-43	Raw Hides, Skins, Leather & Furs	-1.021	5.609	31
44-49	Wood & Wood products	0.578	9.423	31
50-63	Textiles	15.36	13.84	458
64-67	Footwear / Headgear	3.189	26.784	30
68-71	Stone / Glass	3.388	33.419	74
72-83	Metals	2.216	3.766	234
84-85	Machinery / Electrical	-1.655	4.607	418
86-89	Transportation	-9.612	6.021	66
90-97	Miscellaneous	1.253	3.382	195

Note: The table reports the Weighted Average Treatment Effects on the Treated (WATET) exports to Canada after a placebo test, considering the same definitions of treatment but in the period September 2012-August 2015. TET^{*}_{pdt}, are weighted for the relevance each product had in the year before the treaty signature to obtain the unique $WATET$. The weighted standard deviations are computed as $\sqrt{\frac{\sum_{i=1}^{N} s_{pdt}\left(TET^{*}_{pdt}-WATET\right)^2}{(\mathcal{L}-1)\backslash\mathcal{L}\sum_{i=1}^{N} s_{pdt}}}$, where $\mathcal{L}$ is the number of counterfactuals in the trade matrix for Canada. *, **, *** stand, respectively, for $p < 0.05$, $p < 0.01$, $p < 0.001$.

Table A5: Changing alternative destinations in the trade matrix

Model	WATET (1)	weighted std. dev. (2)	N. products (3)
Baseline	1.278***	0.524	2,165
Number of exporters	1,217***	0.423	2,165
Import structure similarity	1.006***	0.431	2,167
Import market size	0.939***	0.429	2,165

Note: The table reports the Weighted Average Treatment Effects on the Treated (WATET) exports to Canada after changing the set of alternative destinations in the trade matrix. TET^{*}_{pdt} are weighted for the relevance each product had in the year before the treaty signature to obtain the unique $WATET$. The weighted standard deviations are computed as $\sqrt{\frac{\sum_{i=1}^{N} s_{pdt}\left(TET^{*}_{pdt}-WATET\right)^2}{(\mathcal{L}-1)\backslash\mathcal{L}\sum_{i=1}^{N} s_{pdt}}}$, where $\mathcal{L}$ is the number of counterfactuals in the trade matrix for Canada. *, **, *** stand, respectively, for $p < 0.05$, $p < 0.01$, $p < 0.001$.

Table A6: CETA and alternative destinations - general equilibrium trade effects - Robustness checks

Dependent variable TET_{pdt}	(1)	(2)	(3)
$TET_{CA,pt}$	-0.927*	-1.951	-1.740*
	(0.332)	(1.250)	(0.648)
$Value_{pdt-1}$	1.755***	1.661***	1.381***
	(0.189)	(0.233)	(0.276)
Constant	-5,379,224.0**	-7,379,045.8*	-4,700,216.0*
	(1,534,383.2)	(3,395,757.2)	(1,947,475.7)
N. obs.	32,505	32,505	32,505
R squared	0.773	0.693	0.602
Clusters by country	Yes	Yes	Yes
Clusters by product class	Yes	Yes	Yes
Model	Number of exporters	Import structure similarity	Import market size

Note: The Table shows results after a linear regression model whose dependent variable includes the treatment effects on the treated in monetary values, TET_{pdt}, where destination d is different from Canada. Each column corresponds to a different set of destinations, as reported in Table A8. The main regressor of interest is the vector of treatment effects on the treated in monetary values, TET_{pdt}, where destination d is instead Canada. The unique control variable is the value of the product p export flow in destination d different from Canada in the period before the CETA, $t-1$. Errors are double-clustered by country and product class. **, *** stand, respectively, for $p < 0.05$, $p < 0.01$, $p < 0.001$.

Table A7: Prediction accuracy at the product level intensive margin - Robustness checks

Model	min RMSE	$\overline{Y}$	SI	NRMSE
Baseline	7.12126	7,060,711	0.000100858	0.00027172
No fixed effects	7.328702	7,060,711	0.000103796	0.00027963
Number of exporters	8.322443	7,037,844	0.000118253	0.00034071
Import structure similarity	9.581219	7,204,660	0.000132986	0.00049488
Import market size	11.518196	7,041,990	0.000163565	0.00053770

Note: The table reports the statistics of the prediction accuracy that we obtain when we train the model while removing the Fixed Effects, or on matrices where we used different matrix structure strategies.

Appendix B: Difference-in-difference

We consider the simple difference-in-difference as a conventional empirical method for benchmarking against our preferred empirical strategy. Following our definitions, a treated product is a product that is enlisted in the CETA, while a treated firm is a firm that exports to Canada at least one product under CETA. Basic formulations are, for the intensive margins:

$$Y_{ut} = c_u + \gamma_t + \beta_D \cdot D_{ut} + \epsilon_{ut} \tag{9}$$

the extensive margin for products:

Table A8: Choice of destinations using different selection criteria

Selection Criterion	Individual Destinations	Aggregates
Baseline	Belgium, Canada, Switzerland, China, Germany, Spain, the United Kingdom, Italy, The Netherlands, Poland, the United States of America	Africa, Americas, Asia, Europe, Oceania
Number of exporters	Belgium, Canada, Switzerland, China, Germany, Spain, the United Kingdom, Italy, Japan, Morocco, the United States of America	Africa, Americas, Asia, Europe, Oceania
Import structure similarity	Austria, Australia, Canada, Germany, Spain, Finland, United Kingdom, New Zealand, Poland, Sweden, The United States of America	Africa, Americas, Asia, Europe, Oceania
Import market size	China, Germany, the United Kingdom, Hong Kong, India, Italy, Japan, Korea, the Netherlands, the United States of America	Africa, Americas, Asia, Europe, Oceania

Note: The table reports, for each destination selection criterion, the list of partner countries included in the trade matrix.

$$Pr(Q_{pt} = 1 | X_{pt} = 1) = c_u + \gamma_t + \beta_D \cdot D_{ut} + \epsilon_{ut} \tag{10}$$

and the extensive margin for firms:

$$Q_{it} = \exp(c_i + \gamma_t + \beta_D D_{it} + \epsilon_{it}) \tag{11}$$

where Y_{ut} represents the total exports of the u-th unit of observation where $u = (p, i)$ is either a p-th product or an i-th-firm observed at time t in Canada. Product fixed effects, c_p, and time fixed effects, γ_t, are included. The binary variable D_{pt} is the treatment indicator, while the error term ϵ_{pt} captures stochastic variation. In eq. 10, we examine the impact of CETA on the product's extensive margin of trade with either a linear probability model (LPM) or a logit, whose dependent variable, Q_{pt} is equal to one if the product was exported and zero otherwise. In eq. 11, instead, we study the impact of CETA on the firms' extensive margin with either a simple OLS or a Pseudo-Poisson, where Q_{it} is the number of products exported in Canada by a firm i at time t.

Table B1: Difference-in difference for products and firms

	Product-level			Firm-level		
	Intensive Margin	Extensive Margin		Intensive Margin	Extensive Margin	
	(OLS)	(LPM)	(Logit)	(OLS)	(OLS)	(Poisson)
	Y_{pt}	$P(Q_{pt}=1)$	$OR(Q_{pt}=1)$	Y_{it}	$Q_{it,CA}$	$Q_{it,CA}$
	(1)	(2)	(3)	(4)	(5)	(6)
ATT	91.63	0.017	1.244	-32.03	0.324***	0.0842***
	(115.1)	(0.010)	(0.161)	(48.52)	(0.075)	(0.028)
Year fixed effects:						
t-4	12.95	0.005	1.067	-5.452	-0.16	-0.008
	(46.53)	(0.005)	(0.085)	(16.65)	(0.030)	(0.010)
t-3	79.97	0.016**	1.228**	21.00	0.127***	0.040*
	(58.09)	(0.005)	(0.085)	(16.65)	(0.049)	(0.016)
t-2	15.83	0.014*	1.201*	-10.05	0.245 ***	0.077***
	(62.16)	(0.006)	(0.087)	(25.52)	(0.061)	(0.020)
t-1	89.88	0.033***	1.526***	18.91	0.420***	0.133***
	(62.61)	(0.006)	(0.113)	(22.83)	(0.070)	(0.021)
t	82.20	0.015	1.219	76.37	0.232***	0.085*
	(128.9)	(0.010)	(0.150)	(50.76)	(0.051)	(0.023)
constant	1,063.6***	0.550***		291.1***	2.407***	
	(45.67)	(0.004)		(16.54)	(0.043)	
Product fixed effect	YES	YES	YES	NO	NO	NO
Firm fixed effect	NO	NO	NO	YES	YES	YES
N. obs.	15,763	31,236	10,980	53,338	53,338	45,729

Note: We report product-level results in columns 1-3. Column (1) reports results on the intensive margin expressed in thousands of euros. Columns (2) and (3) report results on the extensive margin either computed using a Linear Probability model (LPM) or a logit (Logit). We report firm-level results in columns 4-6. Column (4) reports the results on the intensive margin expressed in thousands of euros. Column (5) reports the results on the extensive margin computed using a Linear Model, while column (6) reports results after using a Pseudo-Poisson Model. Robust standard errors in parentheses. *, **, *** stand, respectively, for $p < 0.05$, $p < 0.01$, $p < 0.001$.

Appendix C: Prediction accuracy

Different metrics are used to evaluate the prediction accuracy of machine learning algorithms. Briefly, prediction accuracy metrics compare the classes predicted by the algorithm with the actual ones.

In the case of continuous outcomes, we can use the following measures:

- **Root-Mean-Square Error (RMSE)**, which is computed as

$$RMSE = \sqrt{\sum_{i=1}^{NRD} (\hat{y}_{ird} - y_{ird})^2 / NRD} \tag{12}$$

- **Scatter Index (SI)**, computed as

$$SI = RMSE/\overline{y}_{ird} * 100 \tag{13}$$

  It gives the percentage of expected error for the parameter of interest

- **Normalised Root-Mean-Square Error (NRMSE)**, computed as

$$NRMSE = RMSE/(Q3 - y_{min}) * 100 \tag{14}$$

  it relates the RMSE to the observed range of the variable, thus allowing comparisons with other models