Journal of International Economics 157 (2025) 104145

Contents lists available at ScienceDirect

Journal of International Economics

journal homepage: www.elsevier.com/locate/jie

Full length articles

The tails of gravity: Using expectiles to quantify the trade-margins effects of economic integration agreements

Jeffrey H. Bergstrand [a,b,c,d], Matthew W. Clance [e], J.M.C. Santos Silva [f,*]

[a] Department of Finance, Mendoza College of Business, University of Notre Dame, United States of America
[b] Department of Economics, College of Arts and Letters, University of Notre Dame, United States of America
[c] Kellogg Institute for International Studies, Keough School of Global Affairs, University of Notre Dame, United States of America
[d] CESifo Munich, Germany
[e] Department of Economics, University of Pretoria, Hatfield 0028, South Africa
[f] School of Social Sciences, University of Surrey, Guildford, GU2 7XH, United Kingdom

ARTICLE INFO

Dataset link: The Tails of Gravity (Original data)

JEL classification:
F1
F13
F14
F15
C21
C51

Keywords:
Expectile regressions
Extensive margin of trade
Gravity equation
International trade

ABSTRACT

Although there is evidence suggesting that the effects of trade liberalizations likely vary across the distribution of trade flows, trade economists have focused almost entirely on *conditional mean* estimates of their trade elasticities. We propose the novel use of Poisson-based *expectile regressions* to estimate the heterogeneous effects of trade liberalizations across the entire conditional distribution. Like standard Poisson regression, this method does not need the dependent variable to be logged, accommodates a mass of observations at zero, and is easy to implement, allowing the estimation of gravity equations with the standard three-way fixed effects specification. Using the proposed estimator, we find systematic evidence that trade liberalizations have larger effects at the lower tail of the conditional distribution. We then use the proposed method to investigate the causes of this heterogeneity, and our results suggest that the success of trade liberalizations strongly depends on potential for expansions along the extensive margin.

1. Introduction and motivation

More than sixty years after its introduction by Tinbergen (1962), the gravity equation for trade is now considered “one of the great success stories of recent research on international trade” (Carrère et al., 2020) and, owing to the contributions of Anderson (1979), Helpman and Krugman (1985), Bergstrand, (1985, 1989, and 1990), Anderson and van Wincoop (2003), and others, it has a solid theoretical grounding.[1]

One of the most frequent uses of the gravity equation has been for estimating the partial equilibrium effects of economic integration agreements (EIAs), which can then be used to evaluate the welfare consequences of the policies. This is illustrated, for example, by the recent literature evaluating the consequences of Brexit (see HM Treasury, 2016, Brakman et al., 2018, Dhingra et al., 2017, Felbermayr et al., 2017, Gudgin et al., 2017, and Oberhofer and Pfaffermayr, 2021) and by its use by the U.S. International

* Corresponding author.
E-mail addresses: bergstrand.1@nd.edu (J.H. Bergstrand), matthew.clance@up.ac.za (M.W. Clance), jmcss@surrey.ac.uk (J.M.C. Santos Silva).

[1] Carrère et al. (2020) provide an excellent and up-to-date discussion of the influence of the gravity equation and of its role as the “workhorse” for explaining empirically determinants of international trade flows.

https://doi.org/10.1016/j.jinteco.2025.104145
Received 11 June 2024; Received in revised form 28 July 2025; Accepted 1 August 2025
Available online 13 August 2025
0022-1996/

Trade Commission to analyze the welfare effects of U.S. free trade agreements (see United States International Trade Commission, 2016 and 2021).

Most of these studies use a three-way fixed effects specification of the type introduced by Baier and Bergstrand (2007) to estimate the coefficient of a single EIA dummy on the conditional mean of bilateral trade flows. That is, most studies of the effects of trade liberalizations do not account for their possible heterogeneous effects. However, it is well documented that trade agreements can have widely heterogeneous effects; see, e.g., Baier et al. (2015), Egger and Nigai (2015), Baier et al. (2018, 2019).

The reasons for this heterogeneity vary. Naturally, some agreements are "shallow" while others are "deep". For example, using detailed data on the provisions contained in different agreements, Breinlich et al. (2021) emphasize that EIAs are themselves very heterogeneous and papers such as Baier et al. (2014), Egger and Nigai (2015), Dhingra et al. (2018), and Baier et al. (2019) provide evidence that different types of EIAs can have very different effects, and that the same agreement can have different effects on different countries.

More recently, a different source of heterogeneity in the trade effects of an EIA has surfaced. Specifically, a literature rationalizing the likely sensitivity of the trade elasticity to the *level* of bilateral trade has developed. One of the earliest studies in this literature, Novy (2013), demonstrated using transcendental logarithmic (translog) preferences that the trade elasticity was lower the larger the import share of the good; specifically, the model predicted a larger intensive-margin effect at lower levels of imports *per good* (as a share of expenditures). Also, in a presidential address to the Royal Economic Society, Peter Neary pointed out that, under the assumption of additively separable preferences, the elasticity of bilateral trade with respect to *ad valorem* trade costs is sensitive to the *level* of bilateral trade. Furthermore, assuming subconvexity, (the absolute value of) this elasticity declines with higher levels of trade; see Carrère et al. (2020). This simple and elegant model implies that, as in Novy (2013), an EIA will have an effect that declines with increases in the level of trade flows.

The work of Novy (2013) and Carrère et al. (2020) relies on demand-side and intensive-margin arguments to obtain heterogeneous trade elasticities. However, in the context of a typical Melitz-type heterogeneous firms model with CES demand and intensive and extensive margins, Bas et al. (2017) demonstrated that deviations from the Pareto assumption on productivities implies that the pair-specific aggregate trade elasticity is a function of the (constant) intensive-margin elasticity and a *weighted* extensive-margin elasticity.[2] In their model, the weight is determined by the dispersion in productivities. In particular, in markets where most potential exporters are active (due to relatively low cutoff productivity), the extensive margin change from a fall in trade costs should be small. By contrast, in markets where few exporters are active (due to relatively high cutoff productivity), the extensive margin change from a trade liberalization should be large.[3]

As noted earlier, empirical studies on the impact of EIAs have mostly estimated a single trade elasticity at the conditional mean. However, the sort of heterogeneity identified by Novy (2013), Carrère et al. (2020), and Bas et al. (2017), suggests that EIAs will impact the conditional distribution of trade in complex ways, affecting not only the mean but also other features of distribution, such as its dispersion. Therefore, the elasticity at the mean provides a very limited view of the effect of EIAs, and a much richer picture can be obtained by estimating the trade effects of EIAs along the *entire conditional distribution* of trade. To date, and as described in the next section, the few studies examining how the partial equilibrium effects of EIAs on bilateral trade flows vary across the conditional distribution have used *quantile* regressions; see, e.g., Figueiredo et al. (2014) and Bergstrand and Clance (2025). However, as noted by Santos Silva and Tenreyro (2006, fn. 6, p. 643), the fact that trade data has a substantial percentage of observations equal to zero reduces the appeal of quantile regression in this context.[4]

In this paper, we introduce a novel method to estimate the potentially heterogeneous effects of EIAs across the conditional distribution of trade flows that avoids most of the problems quantile regression faces in this context.[5] Specifically, we propose the use of the asymmetric Poisson maximum likelihood estimator introduced by Efron (1992), to estimate expectile regressions that identify the heterogeneous effects of EIAs across the conditional distribution of trade flows. In contrast to quantile regression, but similar to the Poisson pseudo maximum likelihood (PPML) estimator of the conditional mean (see Gourieroux et al., 1984, and Santos Silva and Tenreyro, 2006), Poisson-based expectile regression readily accommodates zeros in the data without the need for multiple estimation stages. Additionally, the expectiles estimator is easy to implement, allowing the estimation of gravity equations with the standard three-way fixed effects specification, and the estimated parameters are easy to interpret. This is the first paper in the international trade literature to use expectile regression, and one of the very first papers at all to estimate expectiles using Efron's (1992) asymmetric Poisson maximum likelihood estimator.

Because expectile regression is not as popular as most other forms of regression, it is useful at this stage to give some information about expectiles and to contrast them with quantiles. We start by recalling that, in the linear case, quantile regressions are estimated by minimizing a weighted sum of *absolute* residuals, with weights depending on whether the observation is above or below the estimated quantile, and have median regression as a special case when the same weight is given to all observations; see Koenker and Bassett (1978). In turn, in the linear case, expectile regressions are estimated by minimizing a weighted sum of *squared* residuals, with weights depending on whether the observation is above or below the estimated expectile; see Newey and Powell (1987). Therefore,

[2] A sufficient deviation is the assumption of a truncated Pareto distribution; see Melitz and Redding (2015). Bas et al. (2017) assumed a log-normal productivity distribution, which also implies a variable trade elasticity. We will discuss shortly these papers and other related contributions.

[3] Kehoe and Ruhl (2003, 2013) were among the first to note the role of the extensive margin.

[4] Consequently, Figueiredo et al. (2014) and Bergstrand and Clance (2025) employed multi-step methods designed for censored quantile regression to address the zeros issue.

[5] While we will focus on the effects of EIAs, our methods can be used for estimating the heterogeneous effects of any trade frictions or trade-enhancing policies, such as tariff rates.

linear expectile regression has the usual least squares estimator as a special case when the same weight is given to all observations. Likewise, Efron's (1992) asymmetric Poisson maximum likelihood estimator also gives different weights to observations above or below the estimated expectile, and has the PPML estimator of the mean as a special case when the same weight is given to all observations.

Another important aspect to note is that quantiles are local measures of location in the sense that they depend only on the properties of the distribution around the quantile of interest, and therefore are invariant to perturbations in other regions of the distribution. For trade data, this implies that the lower conditional quantiles are identically zero for a substantial part of the population, and therefore do not depend on the regressors, which complicates the estimation of partial effects. In contrast to quantiles, but like the mean, expectiles are global measures of location that depend on global properties of the distribution, and therefore are sensitive to perturbations in any area of the distribution; see Koenker (2013). In the context of trade data, the fact that conditional expectiles are global measures of location has an important advantage: it implies that all expectiles are strictly positive. This is because they depend on the entire distribution and therefore always depend on the positive observations. Therefore, conditional expectiles always depend on the regressors, and consequently do not suffer from the problems that afflict the estimation of quantile regressions in this context.

Despite these differences, expectiles and quantiles share important characteristics. Crucially, both provide information on the location of different regions of the distribution of a variable and therefore both provide information on how the effects of the regressors vary across the distribution. Expectiles and quantiles also share an important drawback: they both suffer from the incidental parameter problem (see, e.g., Lancaster, 2000, and Weidner and Zylkin, 2021). We briefly consider this issue in Section 3, but further research on this topic is left for future work.

Using expectile regressions, we are able to find systematic evidence of heterogeneity, with partial equilibrium effects of EIAs declining as we move towards the top tail of the conditional distribution, consistent with Novy (2013), Carrère et al. (2020), and Bas et al. (2017). Although our results are not directly comparable with theirs, our findings are also broadly in line with those reported by Bergstrand and Clance (2025), who report that EIAs have weaker effects at the top of the conditional distribution. We then investigate the possible sources of heterogeneity and, building on Kehoe and Ruhl (2003, 2013), and Bas et al. (2017), we consider the role of the *extensive* margin of trade in determining the size of the effects of an EIA. This is the first paper to use a broad sample of bilateral trade flows over an extended period of time to examine the hypothesis that trade elasticities depend on the extensive margin of trade. Consistent with Bas et al. (2017), we find that the trade-enhancing effects of EIAs are larger the lower the initial extensive margin to the foreign market. Moreover, in contrast to Novy (2013) and Carrère et al. (2020), we find *much less support* for trade-enhancing heterogeneity along the intensive margin.

The remainder of this paper is organized as follows. Section 2 provides a brief summary of the literature and theoretical context. In Section 3, we describe our novel econometric methodology for the international trade literature and present the results of a small simulation study illustrating its performance in models with high-dimensional fixed effects. In Section 4, we describe the data used. In Section 5, we provide estimation results. In that section, we show that EIA elasticities are lower at higher conditional expectiles, and then we use our broad sample of country-pairs, annual data, and expectile regressions to provide novel estimates of the role of variable extensive margin elasticities towards explaining variable EIA effects. Section 6 concludes and an Appendix provides some additional empirical results.

2. Theoretical context and related literature

We noted in the introduction that several papers have provided theoretical foundations for the heterogeneous partial equilibrium effects of EIAs across the conditional distribution of trade flows, with all papers suggesting that the partial equilibrium effects of EIAs should diminish at higher levels of trade. One of the first papers that examined the endogeneity of the (trade-cost) "trade elasticity" to the level of bilateral trade is Novy (2013). In contrast to most New Quantitative Trade models which assume CES preferences (see e.g., Baier and Bergstrand, 2001, Eaton and Kortum, 2002, Anderson and van Wincoop, 2003, and Melitz, 2003), Novy (2013) uses translog preferences to motivate a structural gravity model. In Novy's (2013) model, the absolute value of the trade elasticity increases with the exogenous number of origin country goods exported and decreases with the share of destination country income spent on the trade flow from i to j. Using a cross-section of bilateral trade flows, Novy (2013) provides evidence that the absolute values of the intensive margin elasticities are negatively related to the imports of j from i as a share of j's expenditures, as predicted by his model. Chen and Novy (2022) expand upon the thesis in Novy (2013), augmenting the empirical work by using panel data and PPML to estimate the standard three-way fixed effects model. Consistent with the results in Novy (2013), the authors find statistically significant and economically plausible effects of currency unions and economic integration agreements on trade, with trade effects declining with the level of the import share of domestic expenditures divided by the total number of products exported by i, that is, the intensive margin trade flow share.

Carrère et al. (2020) found that CES-based structural gravity models have some anomalous implications, also suggesting exploring alternative preference structures. To relax the CES assumption but not entirely lose tractability, they employed additively separable preferences. Such preferences imply that the marginal utility of each good i in each consuming country j depends on "the amount of it consumed". They show that in this case the trade elasticity is variable, a function of the *variable* elasticity of substitution for the country-pair ij. Assuming further subconvexity, for which they note the existence of micro-econometric evidence, the trade elasticity is *decreasing* with increases in the level of trade. They provide cross-section quantile regression evidence of declining (absolute values of) distance elasticities with increases in trade.[6]

In contrast, based on a modified Melitz (2003) model and motivated by the findings in Kehoe and Ruhl (2003, 2013), Arkolakis (2010, p. 1169) used a supply-side argument to predict "that firms with little previous trade will achieve higher growth when variable trade costs fall". At the core of the Arkolakis (2010) approach is the idea that exporting firms reach individual consumers "rather than the market in its entirety". Drawing on the economics of advertising literature that shows empirically that advertising's effectiveness is subject to diminishing marginal returns, Arkolakis (2010) extends the seminal Melitz (2003) trade model to include increasing marginal market-penetration costs that imply increasing convexity of the marginal cost function in reaching additional consumers "beyond the first". In a Melitz (2003) model, a bilateral trade liberalization allows less productive firms to enter a foreign market, and in Arkolakis's model the elasticity of trade with respect to a given trade liberalization is a negative function of the origin country's firms' productivity levels. Arkolakis (2010) shows, using numerical simulations, that this implies higher growth rates of trade from a given liberalization the lower the initial sales of goods, consistent with the observations in Kehoe and Ruhl (2003, 2013). Despite this unambiguous result for firms with low initial sales, the implications of the model at the *aggregate trade* level are less clear. Arkolakis (2010, p. 1196) points out that, to a *first order approximation*, the elasticity of aggregate exports with respect to trade costs is constant, as in Chaney (2008).

We note that the results of Novy (2013), Carrère et al. (2020), Chen and Novy (2022), and Arkolakis (2010) refer exclusively to the intensive margin. Indeed, in the models considered by Carrère et al. (2020) and Novy (2013), the number of exporters is exogenous, *à la* Armington (1969). Although in Arkolakis's (2010) model the number of exporters is endogenous, his result that the elasticity of exports is higher for firms with lower sales is also an intensive-margin result in the sense that it refers to firms that already export to the destination market (labeled there as either intensive margin or new consumers margin).[7]

Yet, in the wake of the seminal papers of Eaton and Kortum (2002) and Melitz (2003), one of the major findings over the last two decades in the international trade literature is that trade-cost liberalizations increase aggregate bilateral trade flows due to the increase in the *extensive* margin of trade, that is, the *number* of firms exporting from i to j. Studies that have provided empirical support to the importance of the relationship between trade-cost variables and the extensive margin include Hillberry and McDaniel (2002), Kehoe and Ruhl (2003), Crozet and Koenig (2010), Kancs and Hove (2010), Lawless (2010), Dutt et al. (2013), Kehoe and Ruhl (2013), and Persson (2013). Determining quantitatively the relative magnitudes of extensive-margin versus intensive-margin effects of trade-cost changes is important because the former have larger implications for exiting and entering of firms, potentially having greater impacts on transitional unemployment.

Surprisingly, only *one* paper has addressed conceptually and empirically a *variable* extensive margin elasticity, akin to the variable intensive margin literature discussed above. Bas et al. (2017) begin with a standard Melitz model with CES demand but, in contrast to most of the literature, assume a log-normal distribution of productivities. This change to the supply side of the model implies that the variable trade elasticity is driven by the pair-specific extensive margin, because the intensive margin elasticity is a constant.[8] As common to a Melitz model, the trade elasticity is the sum of the intensive and extensive margin elasticities. By abandoning the Pareto assumption, their (weighted) elasticity of the extensive margin is a function of a measure of the dispersion of relative firm productivities that varies across country-pairs. Specifically, if the cutoff productivity is low and the market is thick (or, in their terms, "easy"), the marginal entrant from a trade-cost decline will have little influence on aggregate exports due to a smaller impact on the extensive margin (given a large number of exporters serving that market); consequently, the absolute value of the trade elasticity will be small. Conversely, if the cutoff productivity is high and the market is not "easy", then the marginal entrant can have a large impact on the extensive margin (given a small number of exporters that serve the foreign market) and the absolute value of the trade elasticity will be large.[9]

Bas et al. (2017) complement their theoretical model with empirical work using French and Chinese firm-level data of exports and numbers of exporters to various countries in year 2000. An important aspect of their empirical work to support their theory is the use of numbers of each country's exporters that serve foreign markets (as a share of the total number of exporters of each country); we emphasize this point later in our empirical work. Importantly, their Figure 4(a) strongly confirms that the absolute values of the extensive margin elasticities are indeed variable and are decreasing functions of the share of each country's total exporters that serve a particular foreign market. A limitation of the empirical analysis in Bas et al. (2017) is that it employs only data for two countries in a single year (2000). Furthermore, for estimating their (varying) bilateral aggregate trade elasticities, they employ only the "tetrad" method (see, e.g., Head and Mayer, 2014) to avoid all fixed effects. In Section 5, we use a panel of bilateral aggregate

[6] The evidence provided by Carrère et al. (2020) is obtained using the quantile regression estimator proposed by Machado and Santos Silva (2019) which, however, is not appropriate for this kind of data.

[7] Another relevant paper is Spearot (2013). This paper argues theoretically and demonstrates empirically that a common-sized trade-cost reduction can increase low revenue varieties (in an industry) by more than high revenue varieties. The key is the assumption of varying demand elasticities, which is again a demand-side, intensive margin argument.

[8] In Bas et al. (2017), the intensive margin elasticity is $1-\sigma$, where σ is the elasticity of substitution in consumption.

[9] Although we undergird theoretically the variable extensive margin elasticities we find in this paper with Bas et al. (2017), two other papers have provided theoretical foundations for variable extensive margin elasticities. Melitz and Redding (2015) demonstrated theoretically that any small deviations from Pareto in the firms' productivity distribution – such as a *truncated* Pareto distribution – can introduce a variable extensive margin elasticity. Brooks and Pujolas (2019) introduce a general preference structure (additively separable utility) alongside a general technology (constant returns to scale) and intermediate goods that can be aggregated into final goods. Their Proposition 3 shows that the trade elasticity is variable over bilateral trade costs owing to (i) variable sectoral elasticities, (ii) varying curvature of the utility function, and/or (iii) variable compositions of expenditures by country pairs given a change in trade costs. In the interest of brevity, we refer the reader to those papers for more detailed explanations.

trade flows among many countries over 56 years to examine the role of the extensive margin using the standard three-way fixed effects specification from Baier and Bergstrand (2007).

Traditionally, quantile regression has been the econometric approach used to examine the variable effects of EIAs on trade flows over broad samples of data. However, in this context, estimation of quantile regressions is challenging because, as pointed out by Santos Silva and Tenreyro (2006) and discussed in detail below, the conditional quantiles of trade are identically zero for certain values of the regressors. The work by Figueiredo et al. (2016), Baltagi and Egger (2016), and Carrère et al. (2020) are examples of the application of quantile regression in the context of trade data; however, these studies considered only observations with positive trade flows, and therefore did not confront the major challenge to the estimation of conditional quantiles with trade data.[10] In contrast, Figueiredo et al. (2014) explicitly account for the large number of observations with zero trade flows by using a modification of the three-step method for the estimation of censored quantile regression initially introduced by Chernozhukov and Hong (2002) and further developed in Galvão et al. (2013); this enables them to use the standard three-way fixed effects specification of Baier and Bergstrand (2007). However, to alleviate the challenges raised by the estimation of censored quantile regression with a large number of fixed effects, Figueiredo et al. (2014) restrict the effect of the pair-fixed effects to be constant across quantiles. More recently, Bergstrand and Clance (2025) use a similar estimator, but avoid the problems associated with estimation of quantile regressions with multiple fixed effects by adopting the Chamberlain–Mundlak correlated random effects approach of Abrevaya and Dahl (2008).[11]

While these studies find evidence of heterogeneity, they either ignore the zeros or interpret the results obtained with estimators for censored data as results that are relevant for the distribution of observed trade flows including zeros. Moreover, these studies do not investigate the role that the demand and supply sides play in generating heterogeneity along the intensive and extensive margins.[12] We next present a method to estimate conditional expectiles which, like conditional quantiles, provide information on the heterogeneous effects of trade policies on the conditional distribution of observed trade flows, but avoid most of the issues just raised. We then use the proposed method to document the existence of heterogeneity and to shed light on its origin.

3. Econometric methodology

3.1. Expectile gravity

We start by considering the stochastic version of a standard gravity equation which, in its general form, can be written as:

$$X = \exp\left(Z'\beta\right)\eta, \tag{1}$$

where X denotes the trade flow between two trading partners, Z is a vector of explanatory variables, β is a conformable vector of parameters, and η is a non-negative error term such that $\mathrm{E}\left(\eta|Z\right) = 1$, implying that $\mathrm{E}\left[X|Z\right] = \exp\left(Z'\beta\right)$; see, e.g., Santos Silva and Tenreyro (2006).[13]

In gravity equations such as (1), β describes the effect of the variables in Z on the conditional mean of X. However, as emphasized by Santos Silva and Tenreyro (2006), η is generally heteroskedastic, which means that the regressors in gravity equations will have different effects on different regions of the conditional distribution of trade, as predicted by the models of Novy (2013), Carrère et al. (2020), and Bas et al. (2017). In other words, although η has a constant conditional mean (equal to 1), other aspects of its conditional distribution are likely to depend on Z, implying that the explanatory variables affect the distribution of X in complex ways.[14] Therefore, it is interesting to go beyond the conditional mean and study how the effects of the regressors change across the conditional distribution of trade.

As described in the previous section, the traditional way to account for heterogeneous effects in conditional distributions is to use quantile regression. However, as noted by Santos Silva and Tenreyro (2006), the conditional quantiles of X cannot be given by a smooth function such as a gravity equation because trade data typically has a substantial mass point at zero, implying that some quantiles will be identically zero for certain values of the regressors. In fact, when X is non-negative and has a mass point at 0, the qth conditional quantile of X given Z has the form:

$$Q_X\left(q|Z\right) = \begin{cases} 0 & \text{if } q \leq \Pr\left(X=0|Z\right), \\ Q_X\left(\frac{q-\Pr(X=0|Z)}{1-\Pr(X=0|Z)}\middle| Z, X>0\right) & \text{if } q > \Pr\left(X=0|Z\right), \end{cases} \tag{2}$$

where $q \in (0,1)$; see Machado et al. (2016).[15]

[10] Carrère et al. (2020) exclude the zeros because they are only interested in the intensive margin. However, with aggregate data, dropping the zeros does not necessarily isolate the effect on the extensive margin because changes in the positive flows may result from changes in the intensive or extensive margins.

[11] Machado et al. (2016) suggested a different method to estimate quantile regressions for non-negative data with a mass-point at zero but illustrate the application of their method using a different kind of data. Moreover, their method is computationally demanding and it is difficult to link the estimated parameters to a structural model.

[12] Using disaggregated data, Bergstrand and Clance (2025) provide evidence supporting Arkolakis's (2010) finding that trade liberalizations lead to higher growth rates of trade the lower the initial sales of goods.

[13] Note that, as usual, the model can also be written with an additive error term.

[14] Although we often just refer to heteroskedasticity, the same occurs if higher moments, or other features of the distribution, depend on the conditioning variables.

[15] See also Buchinsky and Hahn (1998) for a related result in the context of censored quantile regression.

Table 1
Sample quantiles and expectiles for trade flows in 2017.

θ	0.01	0.05	0.10	0.25	0.50	0.75	0.90	0.95	0.99
$Q_X(\theta)$	0	0	0	0	0.1	12.8	254.3	1,115.0	12,417.2
$\mathrm{E}_X(\theta)$	29.8	139.0	279.8	785.1	2,201.9	6,140.7	16,987.8	33,688.7	147,773.7

Notes: Trade flows are in millions of U.S. dollars. Expectiles are estimated by minimizing a weighted sum of squared residuals while quantiles are estimated by minimizing a weighted sum of absolute residuals. Expectiles provide information on the location of different regions of the distribution, and depend on the entire distribution. Quantiles also provide information on the location of different regions of the distribution, but depend only on the observations around them. We use θ to index both quantiles and expectiles.

The fact that the conditional quantiles of X have a kink at $q = \Pr(X = 0|Z)$ complicates the estimation and interpretation of conditional quantiles for trade data. As discussed in the previous section, although several approaches to the estimation of conditional quantiles for trade have been considered, often these methods are computationally demanding, especially in models with many fixed effects.[16] More importantly, their results are often not easy to interpret because, as Eq. (2) makes clear, the estimated parameters do not have the traditional interpretation as semi-elasticities. Therefore, it is necessary to compute partial effects to be able to interpret the effect of Z on $Q_X(q|Z)$ and this implies that it is necessary to estimate $\Pr(X = 0|Z)$ and to take into account that Z has no effect on the quantiles when $\Pr(X = 0|Z) \geq q$; see Machado et al. (2016) and Bergstrand and Clance (2025).[17]

An alternative way to learn about the effects of regressors on different regions of a conditional distribution is to estimate conditional expectiles, the measures of location introduced by Newey and Powell (1987); see also Efron (1991 and 1992) and Philipps (2022). Heuristically, expectiles can be understood as the expectation in a modified population where some observations are given additional weight. Specifically, for any $\tau \in (0,1)$, the expectile τ of X, denoted $\mathrm{E}_X(\tau)$, can be interpreted as the expectation of X in a population where values of X above the expectile occur $\tau/(1-\tau)$ times as often as they do in the population of interest; see Breckling and Chambers (1988), Efron (1992), and Philipps (2022).[18] This parallels what happens with quantiles, which can be construed as the *median* in a modified population where values of X above the quantile q occur $q/(1-q)$ times as often as they do in the population of interest. It follows from these definitions that $\mathrm{E}_X(0.5)$ coincides with the mean of X, just like $Q_X(0.5)$ corresponds to the median of X.

In the unconditional case, there is some probability mass below each expectile, and therefore each expectile corresponds to a quantile, and vice-versa. However, except in special cases, there is no correspondence between conditional expectiles and conditional quantiles because the probability mass below a conditional quantile is the same for any value of the regressors, while in general this is not the case for conditional expectiles.[19]

It is exactly because conditional expectiles and conditional quantiles have different characteristics that conditional expectiles are useful in contexts where the estimation of quantile regressions is challenging. In particular, as discussed in the introduction, expectiles (including the mean) are global measures of location that depend on global properties of the distribution; see Koenker (2013). In contrast, quantiles are local measures of location that depend only on the properties of the distribution around the quantile of interest. This difference implies that, for non-negative data with a mass-point at zero, conditional quantiles will be zero for some values of Z, while conditional expectiles have the important advantage of being strictly positive for any value of Z. Being global measures of location, expectiles always depend on the positive observations; consequently, in contrast to what happens with conditional quantiles (see Eq. (2)), expectiles can be smooth functions of the regressors, which greatly facilitates their estimation and the interpretation of the results.

The difference between expectiles and quantiles can be illustrated with the data we use in Section 5. Table 1 displays a range of quantiles and expectiles for the trade flows in the year 2017, the final year in our dataset (please see Section 4 for more details on the data). The flows in this subsample vary between 0 and 18,065,991.4 millions of U.S. dollars, with over 35% of the observations being equal to zero.[20] Therefore, in Table 1, all quantiles below 0.35 are identically zero. The median flow is 0.1 millions of U.S. dollars, and then the quantiles increase relatively slowly. In contrast, and as expected, all the expectiles are positive. The mean flow is 2,201.9 millions of U.S. dollars and the upper expectiles are in the extreme tail of the distribution.

The results in Table 1 clearly show the differences between quantiles and expectiles. Being local measures of location, the lower quantiles are dominated by the mass point at zero. In turn, not even the 99th quantile is affected by the very large observations in the upper tail of the distribution, most of which correspond to intra-national trade flows (which represent just over 0.5% of the 36,481 observations in this subsample). In contrast, expectiles, being global measures of location, are never zero and always reflect, to some extent, the upper tail of the distribution. Despite these differences, Table 1 makes clear that both quantiles and expectiles

16 See Figueiredo et al. (2014), Baltagi and Egger (2016), Machado et al. (2016), Carrère et al. (2020), and Bergstrand and Clance (2025).

17 Galvão et al. (2013) note that their three step method to estimate censored quantiles allows for some misspecification in the first stage parametric binary model. However, when used to estimate quantiles for trade data, as in Bergstrand and Clance (2025), the first step plays a role in the estimation of the partial effects, and therefore the binary model needs to be correctly specified. Nevertheless, Bergstrand and Clance (2025) find that their results are not particularly sensitive to the specification used in the first step.

18 Philipps (2022) presents eight other interpretations of expectiles.

19 In particular, as it is well known, the conditional expectation does not generally correspond to a conditional quantile. Efron (1991 and 1992) suggested a method to link conditional expectiles to percentiles by computing the percentage of observations below the conditional expectiles (see also Newey and Powell, 1987). However, quantiles and expectiles have very different properties and so we will not follow that approach; see Koenker (1992, 1993, and 2013).

20 The largest value, USD 18.1 trillion, corresponds to the U.S. internal trade.

provide information on the location of different regions of the distribution of trade flows. The table also shows that, with this kind of data, expectiles provide more information on the upper tail of the distribution and smooth out the mass point at zero, which is the major source of difficulties when estimating conditional quantiles for trade data, and the reason why in this paper we focus on the estimation of conditional expectiles.

To proceed, we assume that the τ-th conditional expectile of η in (1) has the form

$$\mathrm{E}_{\eta}(\tau|Z) = \exp\left(Z'\delta(\tau)\right), \tag{3}$$

where the exponential functional form is used because all expectiles of η are positive, the parameters are indexed by τ to make clear that they vary across expectiles, and $\mathrm{E}(\eta|Z) = 1$ implies that $\delta(0.5) = 0$. Combining Eq. (3) with Eq. (1), we have that the conditional expectiles of X have the form

$$\mathrm{E}_{X}(\tau|Z) = \exp\left(Z'\beta(\tau)\right), \tag{4}$$

with $\beta(\tau) = \beta + \delta(\tau)$. It is interesting to note that, in the special case where η is independent of Z, only the intercept in $\delta(\tau)$ is non zero and all expectiles of X are proportional to each other; that is, the slopes in $\beta(\tau)$ will be the same for all expectiles.[21]

More generally, the way $\beta(\tau)$ varies with τ is informative about how Z affects the conditional distribution of X, and in particular its dispersion. If the slopes in $\beta(\tau)$ increase with τ, an increase in the value of the regressor increases the distance between expectiles, and therefore increases the dispersion of the conditional distribution of X. Naturally the reverse pattern is observed when the slopes in $\beta(\tau)$ decrease with τ. In that case, an increase in the regressor decreases the distance between expectiles and therefore reduces the dispersion of the conditional distribution of X.

Eq. (4) shows that, under our assumptions, expectiles have the form of a traditional gravity equation, and therefore readily accommodate the mass of observations at zero and the estimated parameters can be interpreted as elasticities or semi-elasticities (of the expectile with respect to the regressors). This contrasts with conditional quantiles that, because of their kink at $q = \Pr(X = 0|Z)$, are more difficult to estimate and interpret.

Although most of the limited research on expectiles has focused on the linear model, the asymmetric Poisson maximum likelihood estimator introduced by Efron (1992) can be used to estimate *exponential conditional expectile* models such as Eq. (4). Like PPML, this estimator was originally intended for count data, but it can be used for any data with exponential conditional expectiles; when applied to data that are not counts, we call Efron's estimator the asymmetric Poisson pseudo maximum likelihood (APPML) estimator.

Following Efron (1992), the APPML estimator of $\beta(\tau)$ based on a sample $\{(X_{ijt}, Z_{ijt})\}$ with $i = 1, \ldots, n$, $j = 1, \ldots, n$, and $t = 1, \ldots, T$, is the solution to moment conditions of the form:

$$\sum_{i=1}^{n}\sum_{j=1}^{n}\sum_{t=1}^{T}\omega_{ijt}\left(X_{ijt} - \exp\left(Z'_{ijt}\hat{\beta}(\tau)\right)\right)Z_{ijt} = 0, \tag{5}$$

with

$$\omega_{ijt} = \left|\tau - \mathbf{1}\left(X_{ijt} < \exp\left(Z'_{ijt}\hat{\beta}(\tau)\right)\right)\right|,$$

where $\mathbf{1}(a)$ is the usual indicator function of the event a.

The first order conditions in Eq. (5) can be seen as a weighted version of the first order conditions of the PPML estimator, with the weights being given by ω_{ijt}. For $\tau = 0.5$, ω_{ijt} is a constant and the first order conditions in Eq. (5) coincide with those of the PPML estimator, which impose the orthogonality between the residuals and the regressors. For $\tau < 0.5$, observations below the estimated expectile get more weight than those above and therefore we estimate expectiles below the conditional mean, with the reverse happening for $\tau > 0.5$. That is, the APPML estimator is simply a PPML estimator in which observations receive different weights depending on whether they are above or below the estimated expectile.[22]

Note that, mirroring what happens with PPML, the validity of the APPML estimator depends only on the functional form of the model being correct. Therefore, as long as Eq. (4) holds, $\hat{\beta}(\tau)$ is a consistent estimator of $\beta(\tau)$ and is asymptotically normally distributed with the usual sandwich covariance matrix; see Efron (1992) for details.[23]

Finally, because the models we are interested in contain several sets of fixed effects, it is important to consider the possible consequences of the well-known incidental parameters problem (see, e.g., Lancaster, 2000). The remainder of this section discusses this important issue and presents some simulation results.

[21] As noted by Santos Silva and Tenreyro (2006), this is the only case where the gravity equation can be estimated in its log-linearized form, and the case we use in the simulations below.

[22] In practice, estimation can be performed by repeatedly estimating the model by PPML, with weights that are updated until convergence. The number of iterations required for the estimation of a single expectile typically varies between 3 and 8; therefore estimation of conditional expectiles is generally between 3 to 8 times slower than PPML. However, when estimating multiple expectiles, earlier results can be used as starting values and the number of iterations needed drops to around 3, substantially speeding up the process. A Stata command, *appmlhdfe* (Clance and Santos Silva, 2025), implementing the estimator is available.

[23] As with quantile regression, it is unlikely that all expectiles will be correctly specified. Therefore, we see the estimated expectiles as providing approximations to the true expectiles.

3.2. The effects of the incidental parameters

We are not aware of any research on the topic, but we see no reason to assume that the APPML estimator will inherit PPML's robustness to the incidental parameters problem (see Weidner and Zylkin, 2021). Therefore, although the APPML estimator is consistent when $(n, T) \to \infty$, we expect it to be inconsistent when T is fixed and only $n \to \infty$. However, because the estimator is easy to implement, it can be used in datasets with large n and T, which will reduce the size of the bias.

To gain some insight into this problem, we performed a small simulation study. In view of the results in Chesher and Peters (1994) and Chesher (1995), which suggest that simulation results can be overly optimistic when regressors have symmetric distributions, in our design all the variables are drawn from asymmetric distributions. Specifically, for $i = 1, \dots, n$, $j = 1, \dots, n$, and $t = 1, \dots, T$, the fixed effects ϕ^1_{it}, ϕ^2_{jt}, and ϕ^3_{ij} are obtained independently as draws from the $\chi^2_{(1)}$ distribution. The regressor of interest is then generated as

$$d_{ijt} = \mathbf{1}\left[\left(\phi^0_{ijt} + \phi^1_{it} + \phi^2_{jt} + \phi^3_{ij}\right) < \mathcal{F}_4^{-1}(0.15)\right],$$

with ϕ^0_{ij} also being drawn independently from the $\chi^2_{(1)}$ distribution and $\mathcal{F}_4^{-1}(\cdot)$ denoting the inverse of the cumulative distribution function of a $\chi^2_{(4)}$ random variable (and therefore d_{ijt} is a dummy variable that equals 1 with probability 0.15). Finally, we generate the outcome x_{ijt} as

$$x_{ijt} = \exp\left(\beta d_{ijt} + 0.4\left(\phi^1_{it} + \phi^2_{jt} + \phi^3_{ij}\right)\right)\eta_{ijt},$$

where η_{ijt} is obtained as independent draws from the $\chi^2_{(2)}$ distribution; the $\chi^2_{(2)}$ is highly asymmetric and has mode at zero, implying that this design does not generate values of x_{ijt} equal to zero, but generates a large mass of observations very close to zero. The crucial feature of this design is that, because the errors are independent of the regressors, all expectiles have slope equal to β, and therefore it is easy to compare the biases of the estimates across different expectiles.

Having generated data for $\beta \in \{0, 0.2\}$, $n \in \{45, 90\}$, and $T \in \{15, 30, 45\}$, we estimated β using a 3-way model with it, jt, and ij fixed effects, for $\tau \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$. The results of these simulations, based on 1000 replicas, are summarized in Table 2, which, for each of the cases considered, displays the average value of the estimates and the coverage rates of 95% confidence intervals computed using misspecification-robust standard errors (White, 1980), which are known to be downward-biased in this context (see Weidner and Zylkin, 2021, and the references therein).

The results in Table 2 have several interesting features. First, we notice that there are essentially no biases when $\beta = 0$. This result is important because it suggests that tests for the null of no effect are not affected by incidental parameter bias. However, as expected, we find some bias when $\beta = 0.2$. These biases are towards zero for high expectiles and away from zero for low expectiles, with no noticeable bias at the mean or close to the mean. Naturally, although substantial for $T = 15$, these biases become much smaller as T grows.

The effect of the incidental parameter on the biases is mirrored by the coverage of the 95% confidence intervals. Indeed, the coverage is always good when $\beta = 0$, but for $\beta = 0.2$ we observe that the coverage is relatively low for the smaller values of T, especially for the more extreme expectiles. However, with $T = 45$, the coverage never drops below 84%.

We also performed simulations with $\beta = -0.2$, whose results we do not report in detail, and again found that the biases are towards zero for high values of τ and away from zero for low values of τ. That is, the biases appear to be symmetric around zero, with no bias at zero. Because the simulations do not use any random variables with symmetrical distributions, this finding cannot be explained by the "mirror-image" effect identified by Chesher and Peters (1994) and Chesher (1995), and therefore it does not appear to be a product of the simulations design.

The results of these simulations suggest that, in our application, the biases caused by the presence of incidental parameters are likely to be very small because, as described in the next section, in our dataset n and T are relatively large. Nevertheless, in future work, it would be interesting to see if bias-correction procedures of the type discussed by Weidner and Zylkin (2021) can be used in this context. For now, one should generally be cautious when interpreting confidence intervals for parameters estimated by APPML, especially in models with many fixed effects.[24]

4. Data

We use data on nominal bilateral trade flows in U.S. dollars from the Center for International Prospective Research and Data (CEPII) Gravity Database, which is described in Conte et al. (2022).[25] The CEPII Gravity Database also includes traditional gravity determinants for each possible bilateral trade pair and we merge it with the NSF-Kellogg Institute Database on Economic Integration Agreements constructed by Jeffrey Bergstrand and Scott Baier, which is available at https://sites.nd.edu/jeffrey-bergstrand/.[26] The period of coverage of the combined data is annual from 1962 to 2017 and the potential number of countries in the sample is 193

[24] Whether the incidental parameter problem affects expectile regression more or less than it affects quantile regression is an interesting topic for future research.

[25] The data was downloaded from CEPII (http://www.cepii.fr/CEPII/en/bdd_modele/presentation.asp?id=8) in May 2023.

[26] July 2021 version.

Table 2
Simulation results.

n	T	Expectile				
		10th	30th	50th	70th	90th
		$\beta = 0$				
45	15	−0.001	0.000	0.000	0.000	0.000
		(94.7)	(93.6)	(93.7)	(93.2)	(88.4)
	30	0.000	0.000	0.001	0.001	0.001
		(95.4)	(96.1)	(95.7)	(94.7)	(92.3)
	45	−0.000	0.000	0.000	0.000	0.000
		(95.7)	(95.3)	(95.3)	(94.7)	(92.4)
90	15	−0.001	−0.000	−0.000	−0.000	0.000
		(94.4)	(95.6)	(95.4)	(94.4)	(90.5)
	30	0.000	0.001	0.001	0.001	0.000
		(94.8)	(95.1)	(94.0)	(92.7)	(90.3)
	45	−0.000	−0.000	−0.000	−0.000	0.000
		(95.1)	(94.8)	(94.8)	(94.4)	(93.0)
		$\beta = 0.2$				
45	15	0.224	0.209	0.203	0.197	0.184
		(88.4)	(92.4)	(93.0)	(93.4)	(82.6)
	30	0.214	0.206	0.202	0.199	0.191
		(90.9)	(94.5)	(95.8)	(94.8)	(88.3)
	45	0.211	0.204	0.201	0.198	0.192
		(92.3)	(94.7)	(95.2)	(94.1)	(89.1)
90	15	0.220	0.207	0.202	0.196	0.185
		(75.4)	(91.2)	(95.2)	(93.4)	(69.1)
	30	0.212	0.205	0.202	0.199	0.192
		(83.1)	(92.1)	(94.0)	(93.3)	(82.1)
	45	0.208	0.203	0.201	0.199	0.194
		(85.9)	(94.1)	(94.9)	(93.3)	(84.5)

Notes: The table displays, for different expectiles, the average estimate of β over 1,000 replicas of the simulation and, in parentheses, the empirical coverage of 95% confidence intervals (in percentage). Estimates for the 50th expectile correspond to the standard PPML estimates for the mean.

(especially in the latter portion of the sample).[27] As in Yotov (2012), we construct intra-national trade by subtracting each country's exports from its GDP, which we use as a proxy of a country's national output.[28] Therefore, the number of uni-directional nominal bilateral trade flow observations for the 56 years for which we have the corresponding EIA information is 1,759,865.

Our focus is on EIAs that are defined as one-way preferential trade agreement (non-reciprocal preferential trade agreement), two-way preferential (though not free) trade agreement, free trade agreement (FTA), and a combined grouping for customs unions, common markets, and economic unions, representing "deep" trade agreements (CUCMECU), due to the relative scarcity of each type. The breakdown of economic integration agreement types in our sample and the summary statistics across the agreement types are presented in Table 3; note the relative scarcity of each type of deep trade agreement.

5. Estimation results

5.1. Baseline results

Building on the methodology in Section 3, we will use APPML to estimate a range of models of the form:

$$X_{ijt} = \exp\left(Z'_{ijt}\beta(\tau) + \sum_{s=1962}^{2017} b(\tau)_s \mathbf{1}(i \neq j \wedge s = t) + \varsigma(\tau)_{it} + \vartheta(\tau)_{jt} + \varrho(\tau)_{ij} \right) \eta(\tau)_{ijt}, \tag{6}$$

where Z_{ijt} is a vector containing variables related to EIAs, $\beta(\tau)$ is a conformable vector of parameters measuring the effects of EIAs on different regions of the conditional distribution, $b(\tau)_s$ are the coefficients on a set of dummies defined by $\mathbf{1}(i \neq j \wedge s = t)$ that allow the difference between intra- and inter-national trade to vary over time (see Bergstrand et al., 2015, and Baier et al., 2019),

[27] The EIA data is a balanced panel which includes all current and past possible trade partners, and it is merged with the CEPII dataset which has an indictor for whether a trade partner exist at time t; we keep all country pairs for which both partners exist at t.

[28] As discussed in Bergstrand et al. (2015) and Bergstrand and Clance (2025), there are different ways of obtaining data on intra-national trade. Campos et al. (2021) explored in detail whether the choice between the different approaches matters to structural gravity estimates and concluded that "the estimates for the partial effect of trade agreements on trade flows are very close to each other" (Campos et al., 2021, p. 5). For small countries, subtracting each country's aggregate exports from its GDP can produce negative values, due to the value-added nature of GDP. In our sample, there were only 135 country-year observations that had negative intra-national trade imputations and we replaced these observations with zeros. These 135 country-year observations spanned only 18 exporters. See Yotov (2012, 2022) and Larch et al. (2025) on the importance of using intra-national trade data.

Table 3
Agreements description.

Integration index	Count	Percent of total	Percent of subtotal	Trade Flow, mean (sd)
(0) No Agreement	1,480,145	84.11	84.11	8.2e+07 (1.9e+09)
(1) One-way PTA	153,459	8.72	92.83	2.2e+08 (2.3e+09)
(2) Two-way PTA	52,968	3.01	95.84	2.5e+08 (3.3e+09)
(3) Free Trade Agreement	46,299	2.63	98.47	1.3e+09 (1.0e+10)
(4) Customs Union	11,164	0.63	99.10	1.1e+09 (3.5e+09)
(5) Common Market	10,466	0.59	99.70	2.9e+09 (7.1e+09)
(6) Economic Union	5,364	0.30	100.00	5.2e+09 (1.4e+10)
Total	1,759,865	–	–	1.7e+08 (2.8e+09)

Notes: Total observations are based upon 193 countries for 56 periods (1962–2017), intra-national trade not included. Note that, in the first column, number in parentheses is the number coded in the data source at https://sites.nd.edu/jeffrey-bergstrand. Column (5) provides the mean and standard deviation of trade flows for each agreement type.

Table 4
Baseline results for selected expectiles.

(1)	(2) 10th	(3) 50th	(4) 90th	(5) 10th–90th
EIA_{ijt}	0.314***	0.198***	0.131***	0.183***
	(0.057)	(0.042)	(0.032)	(0.048)

Note: All models include Exporter-year, Importer-year, and Pair fixed effects. Number of observations is 1,499,735. Estimates for the 50th expectile correspond to the standard PPML estimates for the mean. Clustered standard errors by country-pair are in parentheses, * $p < .10$, ** $p < .05$, *** $p < .01$.

$\varsigma(\tau)_{it}$, $\vartheta(\tau)_{jt}$, and $\varrho(\tau)_{ij}$ are the usual three-way fixed effects that account for multilateral resistance and the possible endogeneity of trade policies, and finally $\eta(\tau)_{ijt}$ is an error term with conditional expectile τ equal to 1.[29]

In this section we do not differentiate between the various types of agreements described in Section 4 (see Table 3), and therefore the models include a single EIA dummy that is equal to 1 if any of the six types of agreements is in place; we present results for different types of agreement in the Appendix. We estimate the models for values of τ from 0.02 to 0.98 in steps of 0.01, and then for τ from 0.99 to 0.999 in steps of 0.001. Estimation for values of τ below 0.02 is difficult because almost all observations that get additional weight have trade flows equal to zero, making it difficult to identify the effect of the regressors. This is also reflected by the widening of the confidence intervals as τ approaches 0.[30] Fig. 1 displays these estimates, together with the corresponding 90% and 95% confidence intervals, and Table 4 presents the results for τ equal to 0.1, 0.5, and 0.9, as well as the difference between the estimates for the 10th and 90th expectiles.

The results for $\tau = 0.5$ correspond to standard PPML estimates of the effect of EIAs on the conditional mean, approximately 0.2 (to which corresponds a partial equilibrium effect of approximately 22%).[31] Results for $\tau < 0.5$ give the effect of EIAs in regions of the distribution below the conditional mean, whereas results for $\tau > 0.5$ give the effect of EIAs in regions above the conditional mean. The results in Fig. 1 and Table 4 clearly show that EIAs have very heterogeneous effects across the conditional distribution of trade flows, which decline monotonically as we move to the upper tail of the distribution, keeping everything else constant (e.g., multilateral resistance terms and pair fixed effects).[32] The final column of Table 4 displays the difference between the estimated affects at the 10th and 90th expectiles and the corresponding standard error.[33] These results show that the heterogeneity is both economically and statistically significant, and suggest that EIAs shift up the conditional distribution of trade, with this effect being particularly pronounced in the lower tail. This latter result implies that EIAs not only increase the expected value of trade but also reduce the variance of its conditional distribution.

Our finding that the effect of EIAs decreases with the level of trade (conditional on the covariates) is in line with the results previously obtained using different quantile regression approaches (see, e.g., Carrère et al., 2020 and Bergstrand and Clance, 2025).[34]

[29] We note that the fixed effects vary by expectile to allow the impact of the variables whose effects are subsumed by the fixed effects to vary across the conditional distribution. This contrasts with the more restrictive approach used by Figueiredo et al. (2014), who assume that the fixed effects have a constant impact across the distribution.

[30] A similar problem would occur with quantiles in the extreme left tail, which will be zero for all or almost all observations.

[31] This closely matches the average (weighted by number of country-pairs and inverse of the variance) of the estimates in Baier et al. (2019); see their Table 1.

[32] It is important to keep in mind that these results refer to the conditional distribution, and that low conditional expectiles do not necessarily correspond to low values of trade. A low conditional expectile corresponds to a low value of trade, compared to what would be expected for such pair in a particular year. *Mutatis mutandis*, the same applies to high expectiles.

[33] The standard error for the difference was computed by bootstrap, using 200 replicas and clustering by pair.

[34] Note, however, that our estimates are not directly comparable to those obtained with quantile regression, both because quantiles and expectiles are different functions of the regressors and have different interpretations, and because we estimate (semi-) elasticities of trade with respect to EIAs, which are generally the parameters of interest to policymakers. By contrast, Carrère et al. (2020) and Bergstrand and Clance (2025) estimate elasticities for subsamples with positive trade flows or partial effects (accounting for censoring at zeros), respectively.

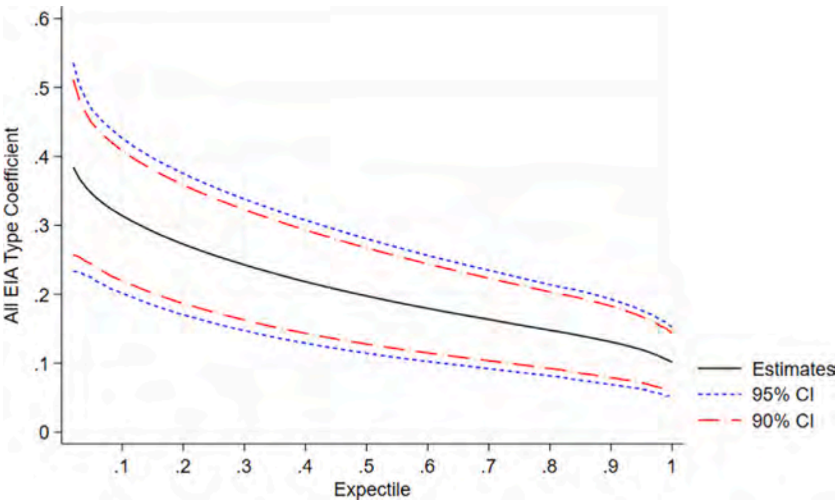

Fig. 1. All trade agreement types.
Note: The binary variable includes all EIA types: (1) One-way PTA, (2) Two-way PTA, (3) Free Trade Agreement, (4) Customs Union, (5) Common Market, and (6) Economic Union.

Moreover, our results are also consistent with the theoretical models of Novy (2013), Carrère et al. (2020), and Bas et al. (2017), all of which predict that the elasticities will be lower for higher levels of trade.

It should be noted, however, that by themselves these results do not allow us to conclude that the effects of EIAs vary along the conditional distribution of trade. In fact, similar results would be obtained if EIAs had a strong and constant effect on the probability of observing positive trade flows and a constant but small effect on the positive flows. Because expectiles are essentially weighted means, the declining effect could simply be an artifact of the diminishing weight given to the zeros as we estimate higher and higher expectiles. However, as we show in the Appendix, a very similar pattern is obtained if the estimation is performed without using the observations for which trade is zero, which suggests that the pattern we observed is not simply the result of EIAs having different impacts on zero and positive flows.

Taken together, these results provide strong evidence that EIAs have heterogeneous effects, but they do not allow us to identify the source of the heterogeneity, and in particular do not allow us to quantify the relative contributions of the extensive and intensive margins to the heterogeneity. We next explore this issue.

5.2. Sources of heterogeneity: The roles of the trade margins

As discussed earlier, Novy (2013), Carrère et al. (2020), and Arkolakis (2010) predict that the partial equilibrium effects of EIAs diminish at higher levels of trade, but these results are exclusively about the intensive margin of trade.[35] However, a key message of the works of Eaton and Kortum (2002), Hillberry and McDaniel (2002), Kehoe and Ruhl (2003), Chaney (2008), Dutt et al. (2013), Kehoe and Ruhl (2013), and others, is that a reduction in trade costs leads to a significant expansion along the extensive margin. The mechanism is that a trade-cost reduction lowers the cutoff productivity level where it becomes profitable for firms to export to the foreign market. Therefore, following a trade liberalization, the number of firms that can profitably export increases, even though the average productivity level of exporting firms declines.

Our earlier results suggest that the effects of EIAs are stronger for the lower tail of the conditional distribution. However, the results in the previous section cannot reveal whether the variable EIA effects are related to variable *intensive* or *extensive* margin effects. In this section, we take theoretical guidance from Bas et al. (2017) and perform an empirical exercise to gauge the extent to which variable extensive margin elasticities explain the heterogeneous effects of EIAs encountered earlier.

As explained earlier, Bas et al. (2017) show that, once the assumption that productivity is (untruncated) Pareto distributed is abandoned, the trade elasticity is a function of a measure of the dispersion of relative firm productivities that varies across country-pairs. Specifically, if the market is thick, the marginal entrant from a trade-cost decline will have little influence on aggregate exports due to a smaller impact on the extensive margin (given a large number of exporters serving that market). Conversely, if the market is not "easy", then the marginal entrant can have a large impact on the extensive margin (given a small number of exporters that serve the foreign market), and the absolute value of the trade elasticity will be large.

[35] More specifically, Arkolakis (2010) focused on, what he termed, the new consumers margin. However, his intensive margin and new consumers margin are effectively an intensive margin, because both of these margins exclude the traditional extensive margin of new firms entering the foreign market.

Table 5
Results for selected expectiles accounting for potential extensive margin expansion.

(1)	(2) 10th	(3) 50th	(4) 90th
EIA_{ijt}	0.657***	0.343***	0.145***
	(0.085)	(0.052)	(0.039)
$\text{EIA}_{ijt} \times \text{EM0}_{ij}$	−0.525***	−0.225**	−0.022
	(0.144)	(0.106)	(0.083)
Total	0.131	0.118*	0.123**
	(0.088)	(0.071)	(0.056)

Note: Total is the sum of the estimates in the top two rows and gives the EIA effect with $\text{EM0}_{ij} = 1$. All models include Exporter-year, Importer-year, and Pair fixed effects. Number of observations is 1,499,735. Estimates for the 50th expectile correspond to the standard PPML estimates for the mean. Clustered standard errors by country-pair are in parentheses, * $p < .10$, ** $p < .05$, *** $p < .01$.

To proceed, we construct a variable similar to that used in Bas et al. (2017) to capture the scope for growth of trade along the extensive margin.[36] Using trade data at the 4-digit level of the Standard International Trade Classification Revision 1, we start by counting the total number of categories of goods with positive exports for each exporter in a given year; this provides a measure of the extensive margin of exports for each country. Then, for each pair and each year, we count the number categories of goods with positive exports from an origin to a destination; this provides a measure of the extensive margin of exports from an exporter to a particular importer. Next, we compute the ratio of the latter to the former, to obtain the share of the number of goods that an origin country exports to a particular partner, as percentage of the total number of goods the origin country exports to the world in a given year.[37] Finally, we create a variable that, for each pair with an EIA, is equal to this ratio in the year *before* the EIA enters into effect, being zero for pairs without an EIA.[38] This variable, labeled EM0, provides for each pair a measure of the potential for growth along the extensive margin when the EIA starts; the lower is EM0, the higher is the potential for extensive margin expansion.

As an illustration, suppose that, in the year before an EIA between countries A, B, and C enters into effect, countries A and B have positive exports in 100 categories and country A has positive exports to country C in all 100 categories, while country B has positive exports to country C in only 50 categories. For the pair (A, C), EM0 is 1, while for the pair (B, C) it is 0.5. Because in country A all sectors that are productive enough to export already export to C, there is little room for trade expansion along the extensive margin. In contrast, many sectors in country B that already export to other destinations may become able to export to C if there is a reduction in the trade costs between these two countries.

The illustration above suggests that the scope for expansion along the extensive margin is likely to be larger for pairs with smaller values of EM0, and therefore we expect that EIAs will have a stronger effect for pairs for which EM0 is smaller. To gauge the importance of the potential for growth along the extensive margin, we modify the specification of the model in (6) to allow the effect of the EIA dummy to vary linearly with EM0. That is, we add to the model the interaction between EM0 and the usual EIA dummy, and expect this interaction to have a negative coefficient.[39]

The results obtained with the re-specified model are presented graphically in Fig. 2, and Table 5 presents results for selected expectiles. Several aspects of these results are noteworthy. We start by noting that the estimates in the first row in Table 5 and the left panel in Fig. 2 correspond to the EIA effect when EM0 = 0, and therefore there is maximum potential for expansion along the extensive margin. These results show that, in this case, EIAs have a strong effect that is larger at lower conditional expectiles. Next, the second row of Table 5 shows that, as expected, the coefficient on the interaction is negative and therefore the effects of EIAs increase with the *potential* for expansion along the extensive margin. Furthermore, this interaction effect is *larger* (in absolute terms) at lower conditional expectiles.[40] Finally, the third row of Table 5 shows the sum of the coefficients in the first two rows, which represents the EIA effect when EM0 = 1. In stark contrast to what we find when EM0 = 0, these results suggest that, when there is little potential for expansion along the extensive margin, the effect of the EIAs is small and stable across the entire conditional distribution. In other words, these results suggest that the intensive margin elasticity is relatively small and fairly constant, ranging between 0.11 and 0.13.

To provide additional insights into how the estimated partial equilibrium effects of EIAs in our sample vary with EM0, Table 6 displays summary statistics for the estimates at different conditional expectiles. For example, for the 90th conditional expectile of the distribution of trade, the minimum of the estimated EIAs effects is 0.131, achieved when EM0 = 1, that is, when EM0 is at its maximum and therefore there is little potential for expansion along the extensive margin.[41]

At the 90th conditional expectile, the estimated coefficient on the interaction is very small (−0.022, from Table 5), so the effect of EIA varies little with EM0. As we move to lower conditional expectiles, the coefficient on the interaction increases in absolute value,

[36] Because they use firm-level data, Bas et al. (2017) define the extensive margin at the firm level. To be able to work with broad sample of bilateral trade flows over an extended period of time, we consider the extensive margin at the sector level, as in Hummels and Klenow (2005).

[37] Alternatively, we could have used as the denominator the total number of categories, but that would not take into account that some countries do not produce some goods and therefore can never export in all 625 categories.

[38] For pairs with EIAs that started before 1962, the first year in our dataset, the variable is equal to the ratio in 1962.

[39] Note that EM0 is time invariant and therefore it cannot be included separately in the model because its role is already captured by the pair fixed effect.

[40] See also the right panel in Fig. 2, noting the y-axis scale; at lower conditional expectiles, the *absolute value* of the interaction effect is larger.

[41] This value is obtained as $\exp(0.145 - 0.022) - 1 = \exp(0.123) - 1$; see Table 5.

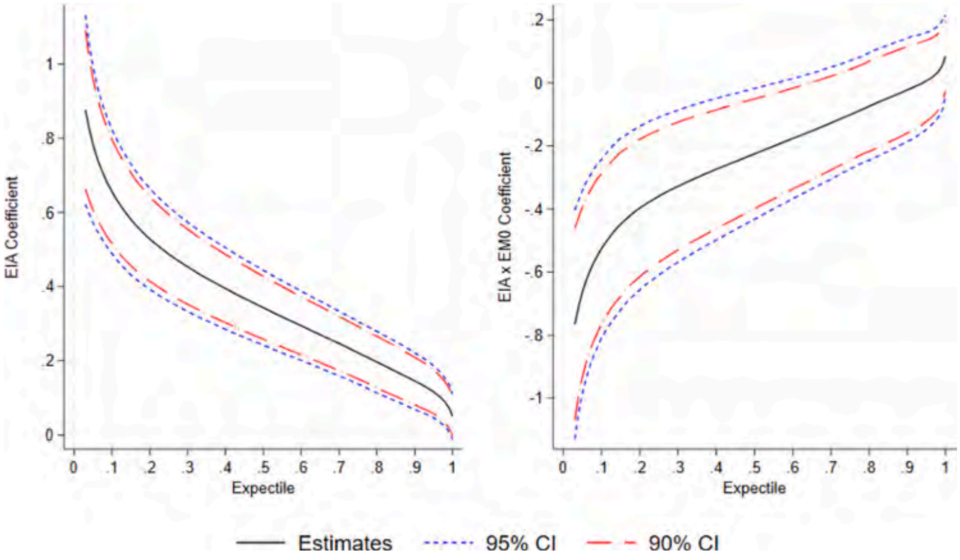

Fig. 2. Extensive margin.
Note: EIA is defined as a binary variable that includes all EIA types: (1) One-way PTA, (2) Two-way PTA, (3) Free Trade Agreement, (4) Customs Union, (5) Common Market, and (6) Economic Union. The left panel is the EIA coefficient without the interaction of EIA with EM0. The right panel is the coefficient estimate for the variable interacting EIA with EM0.

Table 6
Summary statistics of the effects of EIAs at selected expectiles as a function of potential extensive margin expansion.

(1) Expectile	(2) Minimum	(3) Mean	(4) Maximum	(5) Std dev
90th	0.131	0.156	0.156	0.003
50th	0.125	0.401	0.410	0.034
10th	0.140	0.905	0.929	0.099

Notes: The table presents summary statistics of the estimated effects of EIAs at selected expectiles. Because the effect is a decreasing function of EM0, the minimum effect corresponds to the maximum of EM0 (i.e., 1), and vice-versa. Number of observations is 1,769,919. Results for the 50th expectile correspond to the mean.

and therefore the dependence of the effect of EIA on EM0 becomes more pronounced; in Table 6, for the 10th conditional expectile, the estimated effect of EIA varies between 0.140 and 0.929. That is, at the 10th conditional expectile, the partial equilibrium effect of EIAs for country-pairs with large potential for expansion at the extensive margin (0.929) is *more than six times larger* than that of country-pairs with small potential for expansion on the extensive margin (0.140).

Reading Table 6 from top to bottom, we note that the minimum value of the effect (i.e., for EM0 = 1 when there is little scope for expansion along the extensive margin) is essentially constant (varying between 0.125 and 0.140), but there is much more variation as we consider lower values of EM0. For EM0 = 0, when the potential for expansion along the extensive margin is maximal, the effect varies between 0.156 at the 90th conditional expectile and 0.929 at the 10th.

One interpretation of these results is as follows. When EM0 = 1 (which corresponds to the minimum of the EIAs effect in Table 6), there is little scope for expansion along the extensive margin because all sectors in country i that are productive enough to export already export to j. In this case, EIAs affect trade *mostly at the intensive margin*, and the effect is stable across the conditional distribution; this contrasts with the results of Novy (2013) and Carrère et al. (2020) and is more in line with the predictions of the model by Bas et al. (2017). When EM0 is low, there are potentially many sectors in country i that export, but not to j. So, a reduction in trade costs can potentially allow many sectors to start exporting to j. Hence, we see potentially very large effects for lower values of EM0. However, having many sectors that can potentially export does not imply that they will all do so, and hence we see very heterogeneous effects for lower values of EM0: for some pairs a low EM0 will lead to a strong expansion at the extensive margin, but for others it will not, depending on how much of the potential for expansion is realized. This result matches the findings in Bas et al. (2017), shown in their Figure 3(b), that trade elasticities are very dispersed when there are few exporters, but that this dispersion vanishes quickly as the number of exporters grows.

Overall, our results reinforce the findings in Bas et al. (2017) and suggest that heterogeneous responses of the extensive margin as a result of changes in trade costs are the leading cause of the variation of the elasticities of trade across its conditional distribution.

6. Conclusions

In this paper, we suggest the use of Efron's (1992) asymmetric Poisson maximum likelihood estimator to learn about the potential heterogeneous effects of trade costs and trade-enhancing policies on different regions of the conditional distribution of trade flows. The asymmetric Poisson maximum likelihood method estimates conditional expectiles which, like quantile regressions, provide information on the effects of the regressors on different regions of the conditional distribution, but avoid the multi-step censored quantile approach often associated with handling zeros.

Although Efron's (1992) estimator was initially intended to be used with count data, we note that it can be used for any data with exponential expectiles; when applied to data that are not counts, we call Efron's estimator the asymmetric Poisson pseudo maximum likelihood estimator. This estimator is easy to implement, allowing the estimation of large models using large samples, and the estimated parameters have a standard interpretation as (semi-) elasticities.

Using the proposed method, we find that the effects of economic integration agreements are particularly strong in the lower tail of the conditional distribution, a result that is in line with the predictions of Novy (2013), Carrère et al. (2020), Arkolakis (2010), and Bas et al. (2017). In the Appendix we show that this result is robust to changes in the specification of the base model.

We also study the possible causes of this heterogeneity. In particular, we consider the contribution of the extensive margin relative to the intensive margin to explaining the heterogeneity of the effects of economic integration agreements. Our results suggest that heterogeneity in how economic integration agreements affect the extensive margin is a major contributor to the heterogeneous effects of economic integration agreements on the conditional distribution of trade flows. Although we do not claim that our estimates have a structural interpretation, and therefore cannot be used to discriminate between theoretical models, our findings about the role of the extensive margin lend support to the model suggested by Bas et al. (2017), in which the heterogeneity is introduced through the supply-side and affects only the extensive margin, but not to models in which heterogeneity results from departures from the usual CES demand and imply that there is strong elasticity heterogeneity along the intensive margin.

Our results open several interesting avenues for further research. From an empirical point of view, it would be interesting to gather further evidence on how the heterogeneous effects of economic integration along the conditional distribution of trade depend on the extensive margin, and on whether there are also heterogeneous effects along the intensive margin. Also, since the proposed method can be used for estimating the heterogeneous effects of any changes in trade frictions, it would be interesting to see if similar results can be obtained for different trade-enhancing policies, such as reductions of tariff rates. From an econometric perspective, it would be interesting to see if it is possible to use an approach similar to that of Santos Silva and Winkelmann (2026) to study the interpretation of the estimates when the expectiles are misspecified. Additionally, as mentioned before, it would be interesting to study the possibility of using methods such as those discussed by Weidner and Zylkin (2021) to correct the asymptotic bias caused by the incidental parameters problem; this may be particularly important when the samples used do not have a time span as long as the one we used.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We are grateful to the editor Gianmarco Ottaviano and to two anonymous referees for many helpful comments and suggestions. We also thank Piotr Gabrielczak, José Machado, Dennis Novy, Paulo Parente, Maurizio Zanardi, and especially Tom Zylkin, for helpful discussions and advice. We retain sole responsibility for any errors or omissions.

Appendix A

In this appendix we present some additional results that illustrate the application of the proposed APPML estimator and the heterogeneous effects of EIAs.

A.1. Allowing for phasing-in of agreements

Our baseline specification implicitly assumes that the EIAs have a constant effect from the moment they are in place. However, following Baier and Bergstrand (2007), it is standard to account for the fact that EIAs are "phased-in" over time by including lags (and possibly leads) of the EIA dummy. We now consider such a model where we extend our base specification by including two- and four-year lags of the EIA dummy.[42]

Table A.1 presents the estimates obtained with this model for selected expectiles. Besides the usual results, the table also includes the total EIA effect, computed as the sum of the estimates of the parameters associated with the EIA dummy and its two lags. Fig. A.1 presents graphically the estimates of the total EIA effect for a wide range of expectiles.

[42] We experimented with other lags and also leads and the results are qualitatively unchanged.

Table A.1
Results for selected expectiles allowing for phasing-in.

(1)	(2) 10th	(3) 50th	(4) 90th
EIA_{ijt}	0.179***	0.108***	0.069***
	(0.038)	(0.030)	(0.023)
$EIA_{ij,t-2}$	0.075***	0.048***	0.034**
	(0.020)	(0.015)	(0.015)
$EIA_{ij,t-4}$	0.092**	0.063**	0.051**
	(0.041)	(0.029)	(0.022)
Total	0.345***	0.219***	0.153***
	(0.075)	(0.055)	(0.041)

Note: Total is the sum of the estimates in the top three rows. All models include Exporter-year, Importer-year, and Pair fixed effects. Number of observations is 1,387,716. Estimates for the 50th expectile correspond to the standard PPML estimates for the mean. Clustered standard errors by country-pair are in parentheses, * $p < .10$, ** $p < .05$, *** $p < .01$.

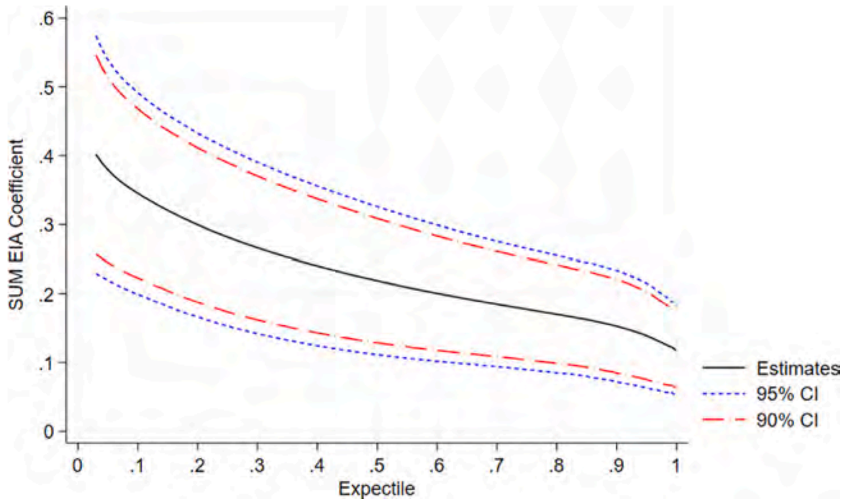

Fig. A.1. Phased-in EIA.
Note: The total EIA effect, defined as the sum of the estimates of the coefficients on EIA_{ijt}, $EIA_{ij,t-2}$, and $EIA_{ij,t-4}$, is graphed across expectiles.

The results with the more flexible specification clearly confirm that EIAs are phased-in over time, with both the second and fourth lags being significant at the standard 5% level for the expectiles whose results are presented in Table A.1. We also note that, in this model, both the total EIA effect and the effects of the three individual dummies decrease as we move from the bottom to the top of the conditional distribution. Nevertheless, the estimates of the total EIA effect allowing for phasing-in are remarkably close to the estimates obtained with the base model. In summary, allowing the effects of EIAs to be phased-in over a period of time reinforces the conclusions obtained with the simpler model used to obtain our baseline results.

A.2. Accounting for differing types of EIAs

So far, we assumed that *all types* of EIAs have the same effect on trade (at the same expectile τ). However, as emphasized, for example, by Breinlich et al. (2021), trade agreements are themselves very heterogeneous in content, with some more shallow than others. Hence, we also expect different types of agreements to have different impacts across the conditional distribution of trade. To account for this *additional source* of heterogeneity, we now consider a richer specification where we include separate dummies for one-way preferential trade agreements (PTAs), two-way PTAs (but not FTAs), free trade agreements, and "deep" agreements (composed of custom unions, common markets, and economic unions).[43] The results obtained with this specification are presented graphically in Figs. A.2, A.3, A.4, and A.5, and the results for selected expectiles are displayed in Table A.2.

The results with the new specification confirm that different types of agreements have very different effects. Furthermore, deep trade agreements have a very strong impact on trade (see, e.g., Dhingra et al., 2018). More importantly, these results reinforce our earlier findings in that we again see that all four types of EIAs considered have effects that decline as we move up the conditional

[43] We include a single dummy (CUCMECU) for "deep" agreements because custom unions, common markets, and economic unions are relatively rare and therefore insufficient variation in these variables makes it difficult to estimate their effects separately.

Table A.2
Results for selected expectiles by type of EIA.

(1)	(2) 10th	(3) 50th	(4) 90th
One-way PTA_{ijt}	0.257***	0.128***	0.072***
	(0.079)	(0.046)	(0.035)
Two-way PTA_{ijt}	0.456***	0.264***	0.138***
	(0.082)	(0.060)	(0.049)
FTA_{ijt}	0.282***	0.198***	0.146***
	(0.065)	(0.053)	(0.042)
$CUCMECU_{ijt}$	0.720***	0.566***	0.452***
	(0.100)	(0.063)	(0.049)

Note: All models include Exporter-year, Importer-year, and Pair fixed effects. Number of observations is 1,499,735. Estimates for the 50th expectile correspond to the standard PPML estimates for the mean. Clustered standard errors by country-pair are in parentheses, * $p < .10$, ** $p < .05$, *** $p < .01$.

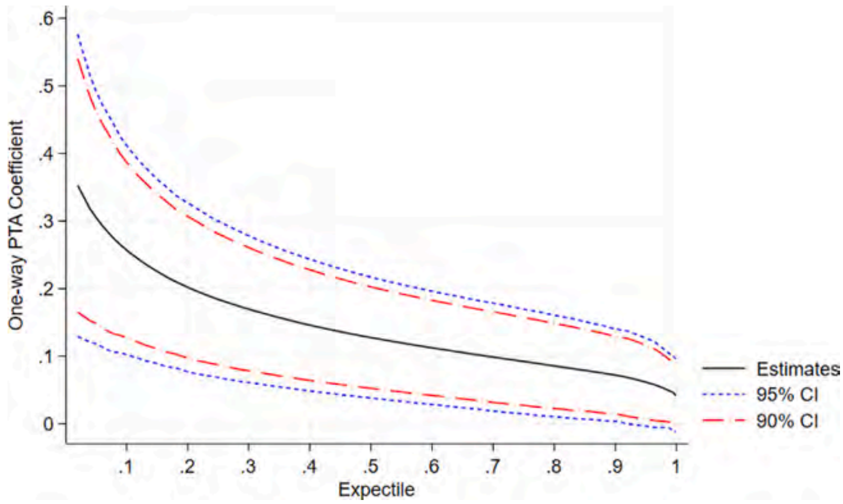

Fig. A.2. One-way preferential trade agreement.
Note: The estimate of the coefficient on the binary variable One-way PTA_{ijt}, or nonreciprocal preferential trade agreement, is graphed across expectiles. The regression also includes Two-way PTA_{ijt}, FTA_{ijt}, and $CUCMECU_{ijt}$.

distribution. In short, our baseline results do not change qualitatively when accounting for the different effects of different types of EIAs.

A.3. Results without zeros

As noted in Section 5, the fact that the estimated effects of EIAs are stronger at the bottom of the conditional distribution than at the top, does not imply that the effects of EIAs vary along the conditional distribution of trade because similar results would be obtained if EIAs had a strong and constant effect on the probability of observing positive trade flows and a constant but small effect on the positive flows.

To shed some light on this issue, we re-estimate our baseline specification excluding the observations for which trade is equal to zero. These results are presented in Table A.3 and Fig. A.6, and have a pattern that is broadly similar to the one obtained using the full sample, suggesting that the pattern we observed is not simply the result of EIAs having different impacts on zero and positive flows.

Bergstrand acknowledges funding from the University of Notre Dame Kellogg Institute for International Studies and the Mendoza College of Business. The authors state that they have no known financial or material interests or personal relationships that influenced the work reported in this paper.

Data availability

The Tails of Gravity (Original data) (Mendeley Data)

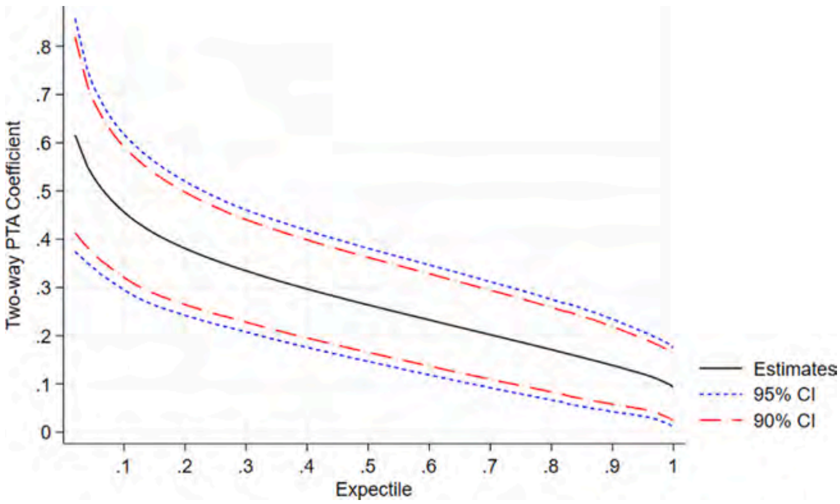

Fig. A.3. Two-way preferential trade agreement.
Note: The estimate of the coefficient on the binary variable Two-way PTA_{ijt}, or two-way preferential trade agreement, is graphed across expectiles. The regression also includes One-way PTA_{ijt}, FTA_{ijt}, and CUCMECU_{ijt}.

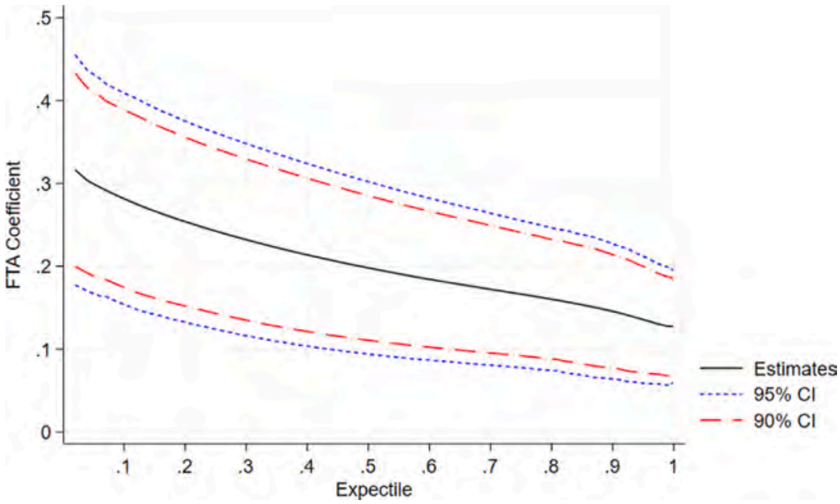

Fig. A.4. Free trade agreement.
Note: The estimate of the coefficient on the binary variable FTA_{ijt}, or free trade agreement, is graphed across expectiles. The regression also includes One-way PTA_{ijt}, Two-way PTA_{ijt}, and CUCMECU_{ijt}.

Table A.3
Baseline results for selected expectiles estimated using only positive values of trade.

(1)	(2) 10th	(3) 50th	(4) 90th
EIA_{ijt}	0.234*** (0.051)	0.152*** (0.039)	0.107*** (0.032)

Note: All models include Exporter-year, Importer-year, and Pair fixed effects. Estimation performed using only positive values of trade. Number of observations is 784,907. Estimates for the 50th expectile correspond to the standard PPML estimates for the mean. Clustered standard errors by country-pair are in parentheses, * $p < .10$, ** $p < .05$, *** $p < .01$.

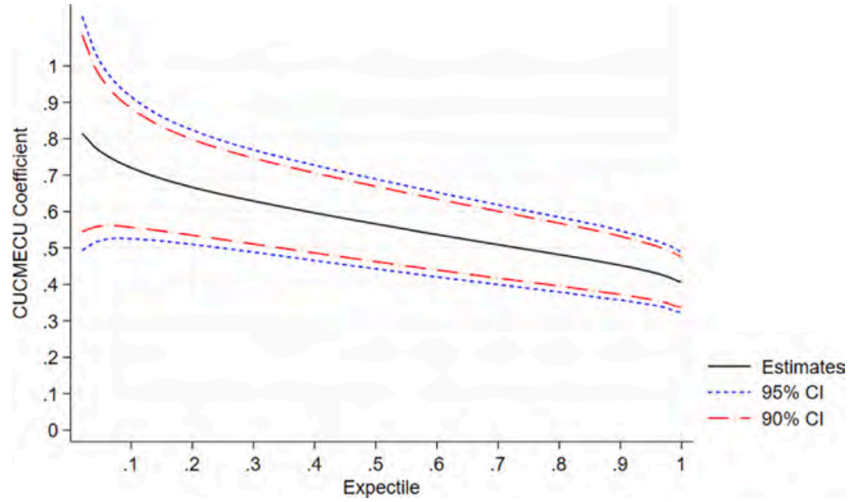

Fig. A.5. CUCMECU.
Note: The estimate of the coefficient on the binary variable $CUCMECU_{ijt}$, defined as deep trade agreements that include customs unions, common markets, and economic unions, is graphed across expectiles. The regression also includes One-way PTA_{ijt}, Two-way PTA_{ijt}, and FTA_{ijt}.

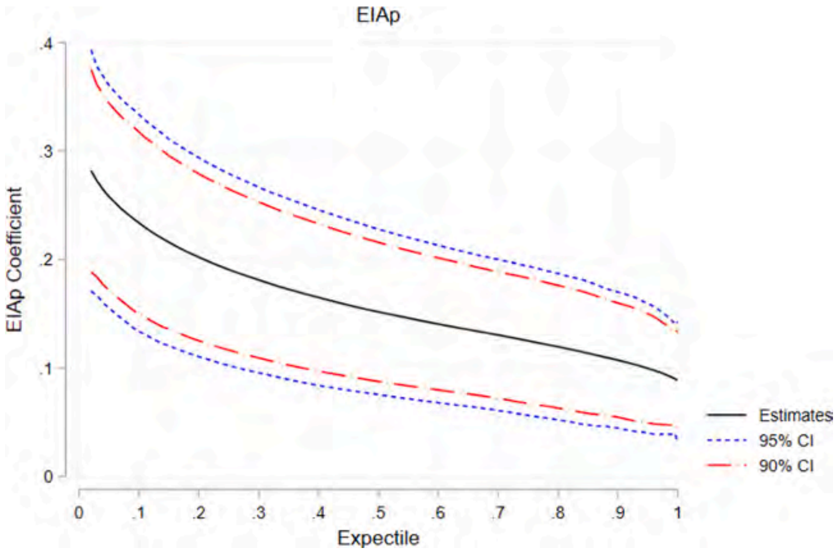

Fig. A.6. All trade agreement types for positive trade.
Note: The binary variable includes all EIA types: (1) One-way PTA, (2) Two-way PTA, (3) Free Trade Agreement, (4) Customs Union, (5) Common Market, and (6) Economic Union. Estimation performed using only positive values of trade.

References

Abrevaya, J., Dahl, C.M., 2008. The effects of birth inputs on birthweight. J. Bus. Econom. Statist. 26 (4), 379–397.
Anderson, J., 1979. A theoretical foundation for the gravity equation. Am. Econ. Rev. 69 (1), 106–116.
Anderson, J.E., van Wincoop, E., 2003. Gravity with gravitas: A solution to the border puzzle. Am. Econ. Rev. 93 (1), 170–192.
Arkolakis, C., 2010. Market penetration costs and the new consumers margin in international trade. J. Political Econ. 118 (6), 1151–1199.
Armington, P., 1969. A theory of demand for products distinguished by place of production. IMF Staff. Pap. 16, 159–176.
Baier, S.L., Bergstrand, J.H., 2001. The growth of world trade: Tariffs, transport costs, and income similarity. J. Int. Econ. 53 (1), 1–27.
Baier, S.L., Bergstrand, J.H., 2007. Do free trade agreements actually increase members' international trade? J. Int. Econ. 71 (1), 72–95.
Baier, S., Bergstrand, J., Clance, M., 2015. Preliminary examination of heterogeneous effects on international trade of economic integration agreements. In: Dur, A., Elsig, M. (Eds.), Trade Cooperation: The Purpose, Design and Effects of Preferential Trade Agreement. Cambridge University Press, Cambridge, UK, pp. 355–373.
Baier, S.L., Bergstrand, J.H., Clance, M.W., 2018. Heterogeneous effects of economic integration agreements. J. Dev. Econ. 135, 587–608.
Baier, S.L., Bergstrand, J.H., Feng, M., 2014. Economic integration agreements and the margins of international trade. J. Int. Econ. 93 (2), 339–350.
Baier, S., Yotov, Y., Zylkin, T., 2019. On the widely differing effects of free trade agreements: Lessons from twenty years of trade integration. J. Int. Econ. 116, 206–226.

Baltagi, B.H., Egger, P.H., 2016. Estimation of structural gravity quantile regression. Empir. Econ. 50, 5–15.

Bas, M., Mayer, T., Thoenig, M., 2017. From micro to macro: Demand, supply, and heterogeneity in the trade elasticity. J. Int. Econ. 108, 1–19.

Bergstrand, J.H., 1985. The gravity equation in international trade: Some microeconomic foundations and empirical evidence. Rev. Econ. Stat. 67 (3), 474–481.

Bergstrand, J.H., 1989. The generalized gravity equation, monopolistic competition, and the factor-proportions theory in international trade. Rev. Econ. Stat. 71 (1), 143–153.

Bergstrand, J., 1990. The Heckscher-Ohlin-Samuelson model, the Linder hypothesis and the determinants of bilateral intra-industry trade. Econ. J. 100 (403), 1216–1229.

Bergstrand, J.H., Clance, M.W., 2025. Quantile gravity: Economic integration agreements, least traded goods, and least developed exporters. Rev. Int. Econ. 33 (4), 951–987.

Bergstrand, J.H., Larch, M., Yotov, Y.V., 2015. Economic integration agreements, border effects, and distance elasticities in the gravity equation. Eur. Econ. Rev. 78, 307–327.

Brakman, S., Garretsen, H., Kohl, T., 2018. Consequences of Brexit and options for a "Global Britain". Pap. Reg. Sci. 97 (1), 55–72.

Breckling, J., Chambers, R., 1988. M-quantiles. Biometrika 75 (4), 761–771.

Breinlich, H., Corradi, V., Rocha, N., Ruta, M., Santos Silva, J.M.C., Zylkin, T., 2021. Machine learning in international trade research: Evaluating the impact of trade agreements.. Policy Research working paper 9629, World Bank, Washington DC.

Brooks, W., Pujolas, P., 2019. Gains from trade with variable trade elasticities. Internat. Econom. Rev. 60 (4), 1619–1646.

Buchinsky, M., Hahn, J., 1998. An alternative estimator for the censored quantile regression model. Econometrica 66 (3), 653–671.

Campos, R., Timini, J., Vidal, E., 2021. Structural gravity and trade agreements: Does the measurement of domestic trade matter? Econom. Lett. (208), 110080.

Carrère, C., Mrázová, M., Neary, J., 2020. Gravity without apology: The science of elasticities, distance, and trade. Econ. J. 130 (628), 880–910.

Chaney, T., 2008. Distorted gravity: The intensive and extensive margins of international trade. Am. Econ. Rev. 98 (4), 1707–1721.

Chen, N., Novy, D., 2022. Gravity and heterogeneous trade cost elasticities. Econ. J. 132 (644), 1349–1377.

Chernozhukov, V., Hong, H., 2002. Three-step censored quantile regression and extramarital affairs. J. Amer. Statist. Assoc. 97 (459), 872–882.

Chesher, A., 1995. A mirror image invariance for M-estimators. Econometrica 63 (1), 207–211.

Chesher, A., Peters, S., 1994. Symmetry, regression design, and sampling distributions. Econometric Theory 10 (1), 116–129.

Clance, M.W., Santos Silva, J.M.C., 2025. APPMLHDFE: Asymmetric Poisson regression with high dimensional fixed effects. Statistical Software Components S459414, Boston College Department of Economics.

Conte, M., Cotterlaz, P., Mayer, T., 2022. The CEPII gravity database. Working Papers 2022–05, CEPII research center.

Crozet, M., Koenig, P., 2010. Structural gravity equations with intensive and extensive margins. Can. J. Econ. 43 (1), 41–62.

Dhingra, S., Freeman, R., Mavroeidi, E., 2018. Beyond tariff reductions: What extra boost to trade from agreement provisions? LSE Centre for Economic Performance Discussion Paper 1532.

Dhingra, S., Huang, H., Ottaviano, G., Pessoa, J., Sampson, T., Reenen, J.V., 2017. The costs and benefits of leaving the EU: Trade effects. Econ. Policy 32 (92), 651–705.

Dutt, P., Mihov, I., Zandt, T.V., 2013. The effect of WTO on the extensive and the intensive margins of trade. J. Int. Econ. 91 (2), 204–219.

Eaton, J., Kortum, S., 2002. Technology, geography and trade. Econometrica 70 (5), 1741–1779.

Efron, B., 1991. Regression percentiles using asymmetric squared error loss. Statist. Sinica 1 (1), 93–125.

Efron, B., 1992. Poisson overdispersion estimates based on the method of asymmetric maximum likelihood. J. Amer. Statist. Assoc. 87 (417), 98–107.

Egger, P., Nigai, S., 2015. Effects of deep versus shallow trade agreements in general equilibrium. In: Dur, A., Elsig, M. (Eds.), Trade Cooperation: The Purpose, Design and Effects of Preferential Trade Agreement. Cambridge University Press, United Kingdom, pp. 374–391.

Felbermayr, G., Groschl, J., Steininger, M., 2017. Britain voted to leave the EU: Brexit through the lens of the new quantitative trade theory. Mimeo: ifo institute.

Figueiredo, E., Lima, L., Schaur, G., 2014. Robust estimation of international trade specifications with heterogeneity in distance and policy effects. Working Paper.

Figueiredo, E., Lima, L., Schaur, G., 2016. The effect of the euro on the bilateral trade distribution. Empir. Econ. 50, 17–29.

Galvão, A., Lamarche, C., Lima, L., 2013. Estimation of censored quantile regression for panel data with fixed effects. J. Amer. Statist. Assoc. 108 (503), 1075–1089.

Gourieroux, C., Monfort, A., Trognon, A., 1984. Pseudo maximum likelihood methods: Applications to Poisson models. Econometrica 52 (3), 701–720.

Gudgin, G., Coutts, K., Gibson, N., Buchanan, J., 2017. The role of gravity models in estimating the economic impact of Brexit. University of Cambridge Centre for Business Research Working Paper No. 490.

Head, K., Mayer, T., 2014. Gravity equations: Workhorse, toolkit, and cookbook. In: Gopinath, G., Helpman, E., Rogoff, K. (Eds.), In: Handbook of Interantional Economics, vol. 4, Elsevier, Amsterdam.

Helpman, E., Krugman, P., 1985. Market structure and foreign trade - increasing returns, imperfect competition and the international economy. The MIT Press, Cambridge, Massachusetts.

Hillberry, R., McDaniel, C., 2002. A decomposition of North American trade growth since NAFTA. United States International Trade Commission Working Paper 2002-12-A.

HM Treasury, 2016. The Long-Term Economic Impact of EU Membership and the Alternatives. HM Treasury, London, UK.

Hummels, D., Klenow, P., 2005. The variety and quality of a nation's exports. Am. Econ. Rev. 95 (3), 704–723.

Kancs, d., Hove, J.V., 2010. Variety gains of trade integration in a heterogeneous firm model. Economics and Econometrics Research Institute Working Paper EERI-RP-2010-34.

Kehoe, T., Ruhl, K., 2003. How important is the new goods margin in international trade. Federal Reserve Bank of Minneapolis Staff Report Number 324.

Kehoe, T., Ruhl, K., 2013. How important is the new goods margin in international trade? J. Political Econ. 121 (2), 358–392.

Koenker, R., 1992. When are expectiles percentiles? Econometric Theory 8 (3), 423–424.

Koenker, R., 1993. When are expectiles percentiles? Econometric Theory 9 (3), 526–527.

Koenker, R., 2013. Discussion: Living beyond our means. Stat. Model. 13 (4), 323–333.

Koenker, R.W., Bassett, G., 1978. Regression quantiles. Econometrica 46 (1), 33–50.

Lancaster, T., 2000. The incidental parameter problem since 1948. J. Econometrics 95 (2), 391–413.

Larch, M., Shikher, S., Yotov, Y., 2025. Estimating gravity equations: Theory implications, econometric developments, and practical recommendations. Working Papers 2025001, Center for Global Policy Analysis, LeBow College of Business, Drexel University.

Lawless, M., 2010. Deconstructing gravity: Trade costs and extensive and intensive margins. Can. J. Econ. 43 (4), 1149–1172.

Machado, J.A.F., Santos Silva, J.M.C., 2019. Quantiles via moments. J. Econometrics 213 (1), 145–173.

Machado, J.A.F., Santos Silva, J.M.C., Wei, K., 2016. Quantiles, corners, and the extensive margin of trade. Eur. Econ. Rev. 89, 73–84.

Melitz, M., 2003. The impact of trade on intra-industry reallocations and aggregate industry productivity. Econometrica 71 (6), 1695–1725.

Melitz, M., Redding, S., 2015. New trade models, new welfare implications. Am. Econ. Rev. 105 (3), 1105–1146.

Newey, W.K., Powell, J.L., 1987. Asymmetric least squares estimation and testing. Econometrica 55 (4), 819–847.

Novy, D., 2013. International trade without CES: Estimating translog gravity. J. Int. Econ. 89 (2), 271–282.

Oberhofer, H., Pfaffermayr, M., 2021. Estimating the trade and welfare effects of Brexit: A panel data structural gravity model. Can. J. Econ. 54 (1), 338–375.

Persson, M., 2013. Trade facilitation and the extensive margin. J. Int. Trade Econ. Dev. 22 (5), 658–693.

Philipps, C., 2022. Interpreting expectiles. Working Papers 2022-01, Department of Economics and Geosciences, US Air Force Academy.
Santos Silva, J.M.C., Tenreyro, S., 2006. The log of gravity. Rev. Econ. Stat. 88 (4), 641–658.
Santos Silva, J.M.C., Winkelmann, R., 2026. Misspecified exponential regressions: Estimation, interpretation, and average marginal effects. Rev. Econ. Stat. (forthcoming).
Spearot, A., 2013. Variable demand elasticities and tariff liberalization. J. Int. Econ. 89 (1), 26–41.
Tinbergen, J., 1962. Shaping the World Economy: Suggestions for an International Economic Policy. The Twentieth Century Fund, New York.
United States International Trade Commission, 2016. Economic Impact of Trade Agreements Implemented under Trade Authorities Procedures, 2016 Report. U.S. International Trade Commission, Washington, DC.
United States International Trade Commission, 2021. Economic impact of trade agreements implemented under trade authorities procedures, 2021 report. U.S. International Trade Commission, Washington, DC.
Weidner, M., Zylkin, T., 2021. Bias and consistency in three-way gravity models. J. Int. Econ. 132, 103513.
White, H., 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. Econometrica 48 (4), 817–838.
Yotov, Y.V., 2012. A simple solution to the distance puzzle in international trade. Econom. Lett. 117 (3), 794–798.
Yotov, Y.V., 2022. On the role of domestic trade flows for estimating the gravity model of trade. Contemp. Econ. Policy 40 (3), 526–540.