

Platforms for Growth?

Multinationals, Trade, and Technology Diffusion

Sheng Cai

John Finlay

Nghiêm Q. Huỳnh*

March 25, 2026

Abstract

We study the role of export platforms—sectors that combine export-oriented production with a strong multinational firm presence—in shaping trade and growth. Empirically, we focus on Vietnam and exploit variation across sectors along two dimensions: first, in exposure to a reform that sharply lowered the cost of exporting to the United States; and second, in pre-reform openness to multinationals. We find that the interaction of export market access and multinational presence generated large increases in exports, sales, and employment. These differential effects were driven by both multinationals and domestically-owned firms, and domestic firms in exposed sectors report more links with multinationals, more technology transfer, and more technology modifications. Motivated by these results, we build a quantitative model that embeds multinational production and technology diffusion within a general equilibrium trade framework. We calibrate the model’s key parameters by targeting our empirical results and quantify how becoming an export platform has contributed to Vietnam’s growth.

Introduction

The foreign affiliates of multinational firms account for roughly one-third of worldwide exports (Cadestin et al. 2018). Much of this activity is concentrated in *export platforms*—countries, sectors, and locations that combine export-oriented production with a strong multinational presence. The *maquiladoras* of Northern Mexico, Bangladesh’s garment sector, and special economic zones throughout China and Southeast Asia are prominent examples. Do these platforms help or hinder the growth of the developing economies that host them?

Policymakers are torn. On the one hand, becoming an export platform allows a developing economy – even one too small or too poor to serve as a destination market – to attract foreign direct investment and benefit from multinational firms’ managerial expertise and technological edge. On the other hand, a natural concern is that multinationals will crowd out local competitors and stifle

*Cai: City University of Hong Kong (email: shengcai@cityu.edu.hk); Finlay: Imperial College London (email: r.finlay@imperial.ac.uk); Huỳnh: University of Oklahoma (email: ngh@ou.edu). We thank seminar participants at the University of Oklahoma, the City University of Hong Kong, International Economics Joint Conference 2025, and FIW 2026 for helpful feedback.

domestic innovation.¹ This debate remains open for two reasons. The first is identification: the emergence of an export platform depends on trade policy and openness to multinational firms, both of which are likely endogenous to other country- or sector-level shocks. The second is aggregation: even with credible sector-level estimates, spillovers and general equilibrium effects may drive a wedge between the partial equilibrium effects of becoming an export platform and its aggregate welfare implications.

We address these challenges in the context of Vietnam, a fast-growing developing country. By combining exogenous changes in trade policy with a shift-share instrument for multinational presence, we provide new empirical evidence on the sector-level effects of becoming an export platform. We use these estimates to calibrate a general equilibrium model that quantifies the role of export platform dynamics. We find that the presence of multinational firms sharply raised Vietnam's gains from a large trade liberalization, but we also find important differences between the partial equilibrium effects we detect with our empirical strategy and the aggregate gains implied by our model.

Vietnam is a paradigmatic export platform, and an ideal setting for our study, for two reasons. First, exports as a share of GDP rose from less than 40% in the mid-1990s to over 80% two decades later, and over a similar period the share of multinationals in total exports rose from 23% to 70% (McCaig, Pavcnik, and Wong 2025; World Bank 2024). The quantitative importance of both international trade and multinational production in Vietnam gives us an excellent opportunity to estimate the interaction of these two forces. Second, as previous work has documented, much of Vietnam's export growth since 2000 has been driven by a plausibly exogenous trade policy reform: the US–Vietnam Bilateral Trade Agreement (BTA) (McCaig 2011; McCaig and Pavcnik 2018). The BTA resulted in large reductions in tariffs on Vietnam's exports to the US, and the size of these reductions varied significantly across manufacturing sectors.

Our empirical exercise asks whether the effects of this trade reform—on exports, revenues, employment, and so on—vary systematically with the pre-reform presence of foreign firms in a sector, which we measure as the share of foreign firms in total revenues. We interpret positive estimates of this interaction as evidence that becoming an export platform amplifies the gains from trade liberalization at the sector level. The challenge, of course, is the endogeneity of foreign firm presence, which we address using our shift-share instrument. Our instrument combines Vietnam's exposure to different multinational origin countries with variation in those countries' sectoral specializations; we show that these forces explain roughly one-fifth of the variation in foreign firm shares across sectors in Vietnam.

Three findings emerge from our empirical analysis. First, we find a large, positive interaction between increases in export market access caused by the BTA and foreign firm presence: a 1% reduction in trade costs raises exports by 10% in a typical sector, but by 19% in a sector with a strong foreign firm presence. Second, we find that this positive interaction is the product of two forces:

¹For example, in a World Bank review Chandra and Kolavalli (2006) write that "...without host-country policies to develop local capabilities, MNC-led exports are likely to remain technologically stagnant, leaving developing countries unable to progress..."

larger direct effects of increased export market access on foreign firms and positive spillovers to domestic producers in export-platform sectors. Third, survey evidence points to an important role for technology diffusion in generating these spillovers for domestic producers. Taken together, our empirical results suggest large gains from combining export market access with openness to multinational firms. Moreover, these gains are not concentrated solely among foreign firms but instead diffuse to domestic producers. Whether these partial equilibrium, sector-level export platform effects translate into aggregate welfare gains is the question we turn to with the model.

We build a quantitative model that embeds multinational production and technology diffusion within a general equilibrium trade framework. Two forces in the model map directly to our empirical results. The first is multinational mobility: because multinational firms can choose to produce in any of many locations, they are more responsive to shocks in a single country than domestic firms. This force can explain why sectors that are more open to multinational firms respond more to reductions in trade costs, but it cannot account for the positive spillovers to domestic firms that we detect in the data. Instead, these spillovers arise from the second force in the model: learning from export platforms. A reduction in trade costs induces multinationals to enter and expand, and because multinationals are highly selected and therefore highly productive, their entry gives domestic producers new opportunities to learn and improve their technologies. These two forces together generate the positive export platform effects we find in the data.

We calibrate the model using data on global trade and multinational flows as well as our empirical results from Vietnam. We then conduct counterfactual exercises to quantify Vietnam's gains from the BTA, which we find are substantial — welfare rises by 1.24% in present value terms. Multinationals play a central role in these gains: if Vietnam were counterfactually closed to foreign firms, the BTA would have raised welfare by a mere 0.22%. Surprisingly, we find that the short-run welfare gains from the BTA overshoot the long-run gains, implying that the gradual process of technology diffusion set in motion by the BTA actually reduces Vietnamese welfare. The explanation for this pattern is that the forces we highlighted above – multinational mobility and learning from export platforms – also imply that multinationals are particularly likely to relocate production away from Vietnam as wages rise in equilibrium, leading to a loss of positive spillovers in the aggregate.

Overall, our empirical results point to gains from the interaction of increased export market access and foreign firm presence: becoming an export platform is a net plus for an individual sector. Our model explains these effects in terms of the greater mobility of multinationals and their ability to generate positive spillovers to domestic firms. But our model also suggests caution in naively extrapolating sector-specific effects into aggregate welfare gains: a positive shock in one sector may well crowd out highly productive multinationals, and with them opportunities for learning, in another sector.

Related literature

A large empirical literature studies spillovers from multinational firms and their implications for host-economy outcomes (see [Harrison and Rodríguez-Clare \(2010\)](#) and [Keller \(2022\)](#) for reviews). Evidence on within-sector spillovers is mixed. In early work, [Aitken and Harrison \(1999\)](#) find negative spillovers to Venezuelan firms, while [Keller and Yeaple \(2009\)](#) detect positive spillovers to American firms. As [Keller \(2022\)](#) points out, beyond the endogeneity of multinational presence, this literature also faces the challenge of disentangling true spillovers between firms—i.e., externalities that might warrant policy intervention—from other equilibrium interactions, such as competition in input or product markets. This paper is the first to estimate how multinational presence interacts with export-market access, by combining policy-induced variation in trade costs with a new instrument for multinational presence. We then isolate the role of spillovers by calibrating a quantitative model in which domestic firms and multinationals interact in a rich input-output structure following [Caliendo and Parro \(2015\)](#).

We also contribute to a quantitative literature that measures the gains from multinational production and international trade ([Ramondo and Rodríguez-Clare 2013](#); [Tintelnot 2017](#); [Arkolakis et al. 2018](#)), and build on [Buera and Oberfield \(2020\)](#) and [Cai and Xiang \(2025\)](#) to incorporate technology diffusion. Our theory emphasizes a new mechanism that explains why export-oriented multinationals are particularly likely to generate positive spillovers: these firms must overcome both the costs of multinational production and the iceberg costs involved in exporting, and so are more selected, and hence more productive, than their counterparts that overcome only one of these obstacles.

Finally, we contribute to a growing literature that explores the causes and consequences of Vietnam’s rapid growth in recent decades. An empirical literature documents the BTA’s effects on poverty ([McCaig 2011](#)), formalization and labor reallocation ([McCaig and Pavcnik 2018](#)), wages and skill premia ([Fukase 2013](#)), labor market distortions ([Hoang, Mitra, and Pham 2025](#)), family formation ([Drozdoff, McCaig, and Pavcnik 2025](#)), and intergenerational mobility ([Mitra, Pham, and Ural Marchand 2026](#)). These studies provide important insights into the partial equilibrium effects of the BTA along a diverse set of margins, while ours is the first to connect such partial equilibrium estimates to a theory of trade and growth and thus quantify the aggregate importance of the BTA, as well as trade and openness to multinationals more broadly, in Vietnam’s development.

The rest of the paper is structured as follows. Section 1 uses the BTA and its interaction with foreign firm presence to estimate export platform effects. Section 2 develops a quantitative model of trade, multinational production, and technology diffusion and Section 3 calibrates the model by targeting our empirical results. Section 4 uses the model to quantify the importance of trade and multinational firms in Vietnam’s growth. Finally Section 5 concludes.

1 Empirical Evidence of Export Platform Effects

1.1 Data

Our empirical analysis relies on four datasets. Two are standard: we take bilateral trade data from UN Comtrade/BACI (1996–2015), and we draw on firm-level information from the historical Orbis database (1996–2019) to construct an instrument for foreign firm presence, discussed in more detail below.

We also use firm-level data from the Vietnam Enterprise Survey (VES), an annual census conducted by the General Statistics Office (GSO). The survey is compulsory for all registered enterprises in the formal sector. Between 2000 and 2015, the VES covers roughly 580,000 firm–year observations and 150,000 unique firms in manufacturing. We classify firms in the VES into three mutually exclusive ownership groups based on initial equity structure. *Foreign firms* are those with any foreign equity stake: in practice 77% of foreign employment in 2000 was concentrated in 100% foreign-owned enterprises. *Domestic firms* are firms which are not foreign and which are less than 50% state owned. Finally, *state firms* are those with majority or full state ownership.

Finally, we use the Vietnam Technology and Competitiveness Survey to measure technology diffusion and innovation. This survey, conducted between 2010 and 2014, is an official VES module, jointly designed by the GSO, the Central Institute for Economic Management, and the Development Economics Research Group at the University of Copenhagen (Newman et al. 2015). The 2010 round sampled approximately 7,900 manufacturing firms (6,270 private and 1,670 foreign), covering 43% of private and 48% of foreign manufacturing employment (McCaig, Pavcnik, and Wong 2025).

1.2 Context

Vietnam is an ideal setting in which to study the interaction between trade liberalization and foreign firms. Over the past three decades, the country has rapidly integrated into the global economy, with exports rising from less than 40% of GDP in the mid-1990s to over 80% by the mid-2010s. Foreign firms played a central role in this transformation. In 2000 they accounted for 23% of total exports, but by 2020 this figure had risen to 70% (McCaig, Pavcnik, and Wong 2025; World Bank 2024). Vietnam is not an outlier: in China the foreign-owned share of exports jumped from 10% to almost 60% between 1990 and 2004 (Blonigen and Ma 2007).

Trade liberalization

What makes Vietnam particularly useful for the present paper is the fact that a large share of its export growth was driven by a plausibly exogenous trade policy reform: the 2001 US–Vietnam Bilateral Trade Agreement, or BTA. Prior to the BTA, Vietnam’s exports to the United States faced the *Column 2* tariff schedule, with average *ad valorem* rates of around 32%. In December 2001, these

tariffs were replaced by the lower Most-Favored-Nation (MFN) rates, averaging about 3%.² Panel (a) of Figure 1 illustrates the magnitude of this change by plotting the distributions of pre- and post-BTA tariffs across 4-digit Standard Industrial Classification (SIC) sectors. The BTA caused the entire distribution of tariffs to shift sharply to the left, but with substantial variation across industries. Since the new schedule was pre-determined and applied to all industries without exception, the distribution of tariff cuts induced by the BTA may reasonably be regarded as exogenous to other industry-level shocks. We denote an industry i 's exposure to the BTA by T_i , defined by

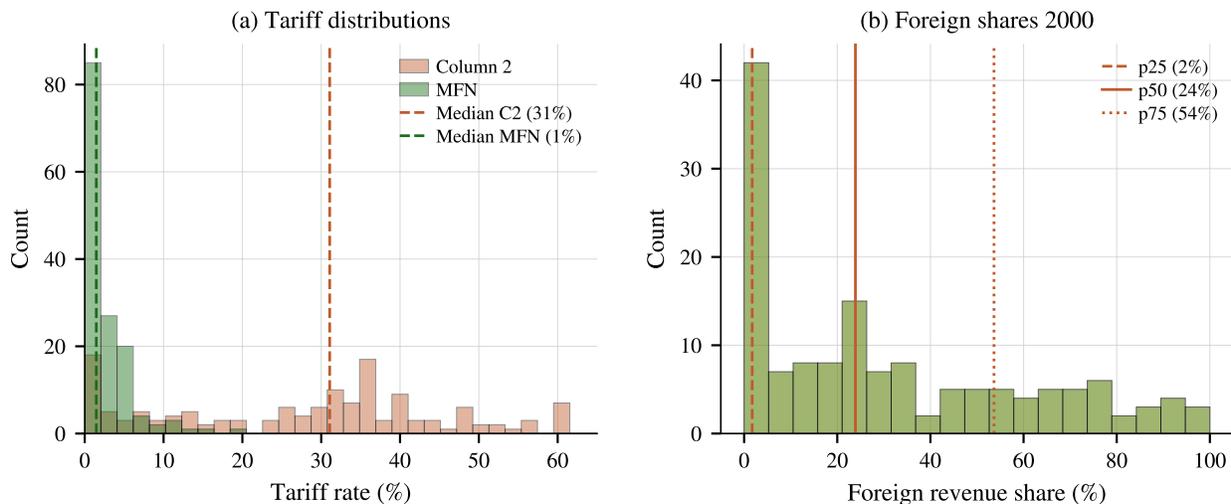
$$T_i = \log(1 + \tau_i^{\text{Column 2}}) - \log(1 + \tau_i^{\text{MFN}}).$$

Notice that T_i is defined so that it is positive for an industry experiencing a tariff reduction under the BTA. We therefore expect its effects on, e.g., Vietnam's exports to the US, to be positive.

Foreign presence

We measure foreign presence using the share of foreign firms in total sector revenue in 2000, and henceforth refer to this as the *foreign share*. We denote this quantity by F_i , and plot its distribution across sectors in Panel (b) of Figure 1. Foreign shares are very dispersed across sectors. The median sector has a foreign share of about 0.24, but the 25th percentile of this distribution is just 0.02, and

Figure 1: The BTA and foreign firm shares



Notes: Panel (a) shows the shift in the U.S. tariff schedule for Vietnamese goods from Column 2 rates to MFN rates in December 2001 using data from [McCaig, Pavcnik, and Wong \(2025\)](#). Panel (b) shows the distribution of foreign firm revenue shares across Vietnamese manufacturing industries in 2000 based on authors' calculations using the VES.

²In contrast, prior to the BTA, Vietnam had already granted MFN tariff status to imports from the United States. As a result, the BTA did not substantially alter Vietnamese tariffs on U.S. exports. Beyond tariff reductions, the agreement required Vietnam to implement customs reforms consistent with WTO standards, liberalize trade in services, strengthen intellectual property rights protection, enhance the investment climate for foreign firms, and increase regulatory transparency ([STAR-Vietnam 2003](#); [McCaig 2011](#)).

at the 75th percentile is 0.54. This heterogeneity in pre-BTA foreign presence will form the basis of our empirical strategy, which asks whether industries with a stronger foreign presence experienced larger or smaller BTA effects.

A natural concern is that if T_i and F_i are highly correlated, we might struggle to estimate their interaction. In fact the correlation between these variables is just 0.187. Appendix Figure B1 shows a scatterplot which makes it clear that there are many sectors with low foreign shares and high BTA exposure, and vice versa. This is the variation we will rely on to estimate their interaction.

In the regressions below, we demean both T_i and F_i so that the direct effects are evaluated at the mean of the other variable. This is purely for ease of interpretation.

1.3 Conceptual Framework

We begin by sketching a simple model of Vietnam's exports to the United States, in order to clarify the mechanisms that might be at work and to motivate the empirical specifications that follow. The model is deliberately stylized: Section 2 will develop a richer quantitative framework that relaxes many of the simplifying assumptions made here.

Setup. Two types of firm operate in Vietnam: foreign-owned (F) and domestic (D). There are many sectors, indexed by i , and all firms sell to a single destination market, the United States. For each firm type $t \in \{F, D\}$ and in each sector, sales take a standard gravity form,

$$X_{it} = \left(\frac{A_{it}}{\tau_i} \right)^{\theta_t}, \quad (1)$$

where A_{it} is a composite productivity term, τ_i is a sector-specific iceberg trade cost, and $\theta_t > 0$ governs the trade elasticity – which we allow to differ across firm types. We model the BTA as a sector-specific shock to trade costs and define $T_i = -d \log \tau_i$.

Direct effects. Suppose first that for both firm types the productivity term A_{it} does not respond to changes in trade costs. Then, following the BTA,

$$d \log X_i = \theta_D T_i + (\theta_F - \theta_D) F_i T_i, \quad (2)$$

where F_i is the foreign share in sector i . This expression suggests a natural regression specification with a direct effect of the trade cost shock and an interaction with the sector-level foreign firm share. The interaction term is positive whenever $\theta_F > \theta_D$, that is, whenever foreign firms are more responsive to falling trade costs than their domestic counterparts. This might be the case if, for example, multinationals are able to adjust to shocks by reallocating production across countries as in [Ramondo and Rodríguez-Clare \(2013\)](#).

Spillovers. It is natural to expect that domestic and foreign firms might interact, either through general equilibrium channels or via productivity spillovers. To capture this, we now allow domestic productivity A_{iD} to depend on the foreign share,

$$A_{iD} = (1 - \lambda) + \lambda F_i, \quad (3)$$

where $\lambda > 0$ implies that foreign presence raises domestic productivity – for instance through technology diffusion or supplier linkages – and $\lambda < 0$ implies crowding out. Taking a first-order approximation around $F_i = 0$, the response of exports by domestic firms becomes

$$d \log X_{iD} = \theta_D T_i + (\theta_F - \theta_D) \theta_D \lambda F_i T_i, \quad (4)$$

while for total exports we obtain

$$d \log X_i = \theta_D T_i + (\theta_F - \theta_D) (1 + \theta_D \lambda) F_i T_i. \quad (5)$$

The interaction coefficient for total exports now captures both a composition effect ($\theta_F - \theta_D$) and any spillovers to domestic producers (λ).

Takeaways. We refer to the positive interaction between trade cost reductions and foreign firm presence as the *export platform effect*. Splitting the sample by firm type allows us to separate the *composition* channel ($\theta_F > \theta_D$) from the *spillover* channel ($\lambda \neq 0$): if positive spillovers are present, we should see positive interaction effects for domestic firms, not just larger direct effects for foreign firms.

1.4 Empirical Strategy

Guided by the conceptual framework above, our main regression specification is

$$Y_i = \alpha + \beta T_i + \gamma F_i + \delta T_i F_i + u_i, \quad (6)$$

where i indexes 4-digit SIC sectors, Y_i is the outcome of interest, and δ , the parameter of interest, captures the export platform effect.

Given that the BTA moved Vietnam from one pre-defined tariff schedule to another, T_i can safely be regarded as exogenous. Instead, the key threat to the identification of δ is that foreign firms may sort into sectors that would anyway respond differentially to the BTA. If foreign firms cluster in inherently dynamic industries that are better able to expand when tariffs fall, we would overestimate δ . If they instead cluster in sectors where domestic competitors are weak and unable to take advantage of lower tariffs, the bias runs in the opposite direction. To address this challenge, we now construct an instrument for foreign presence Z_i , and then use it and its interaction with T_i to estimate (6) by two-stage least squares (2SLS).

Instrumenting for the foreign share

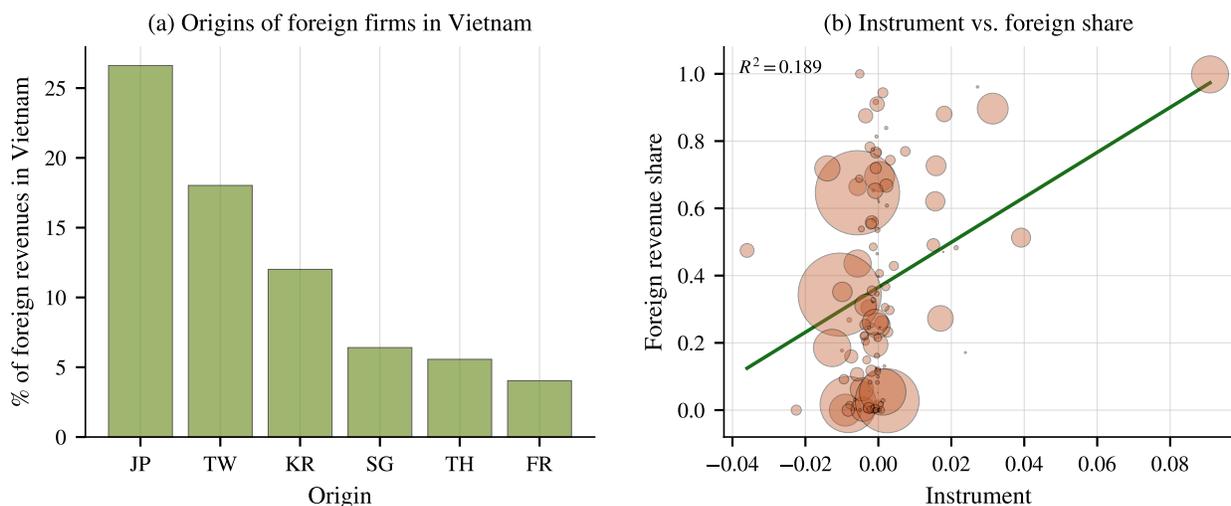
Our instrument builds on two observations. First, multinationals in Vietnam are drawn from a specific set of origin countries. Second, these origin countries tend to produce multinationals in specific sectors. Thus, Vietnam’s overall exposure to certain origin countries allows us to predict variation in the share of foreign firms across sectors.

Panel (a) of Figure 2 illustrates the first point: in 2000, Japan accounted for 26.6% of all foreign firm revenues in Vietnam, Taiwan for 18.0%, and South Korea for 12.0%. We denote these origin revenue shares by w_o , where o indexes the 47 countries that appear in the Orbis data.

To illustrate the second point, consider SIC code 3430, parts and accessories for automobiles. Orbis data shows that Japan derives 3.46% of its total multinational revenues from this sector, far above the cross-country average of 0.32%, and South Korea derives 1.02%. Given Vietnam’s high exposure to these countries, we would predict an unusually high foreign share in this sector – and indeed, the foreign share in 2000 was 48.3%, versus a median of 23.9%. Importantly, this prediction uses no information specific to the automobile parts sector in Vietnam.

Our IV strategy applies this idea systematically. Let S_{io} be the share of multinational revenues derived from sector i in origin country o in total multinational revenues in origin country o , measured using Orbis data between 1996 and 2019. Let \bar{S}_i denote the mean of this share across countries, and define $\hat{S}_{io} \equiv S_{io} - \bar{S}_i$. Intuitively \hat{S}_{io} measures the extent to which country o is unusually likely to generate multinational revenues in sector i . We then define our instrument for

Figure 2: Construction and Validation of the Foreign Presence Instrument



Notes: Panel (a) shows the share of total foreign firm revenues in Vietnam attributed to foreign firms originating in the six largest origin economies: Japan, Taiwan, Korea, Singapore, Thailand, and France. Panel (b) shows a sector-level scatterplot of the instrument defined in (7) against foreign revenue shares calculated using the VES. Marker sizes are proportional to each sector’s share of Vietnam’s exports in 2000 and the solid line shows a weighted least squares regression.

the foreign share in Vietnam in sector i by

$$Z_i = \sum_o w_o \hat{S}_{io}. \quad (7)$$

Panel (b) of Figure 2 shows the instrument Z_i and the foreign share F_i are positively correlated. A regression of F_i on Z_i results in a coefficient of 4.74 with a standard error of 0.98, significant at the 1% level, and an R^2 of 0.189.

These statistics show that Z_i satisfies the instrument relevance criterion. We also require Z_i to be uncorrelated with unobserved determinants of sector-level outcomes. For this to fail, there would need to exist some omitted sectoral characteristic that simultaneously causes a sector to respond differentially to the BTA, raises the share of that sector in total multinational revenues in the countries from which Vietnam is particularly likely to draw multinationals, and yet does not raise the share of that sector in total multinational revenues globally. Omitted characteristics that are common to all countries, like the extent of product differentiation across firms, or specific to a single sector in Vietnam, like industrial policy targeting a favored sector, can therefore be ruled out as potential sources of bias.

1.5 Results

Export platform effects

The first prediction of the conceptual framework is that sectors with a higher foreign share should respond more strongly to the BTA. We test this prediction using the outcome most directly connected to the BTA, Vietnam's exports to the US, measured by log differencing each sector's exports between 2000 and year t , for t between 2002 and 2015. In all the specifications below we include year fixed effects, weight sectors by their share in Vietnam's total exports in 2000, and cluster standard errors by sector.

In Column (1) we estimate the simplest possible specification by OLS, without including the foreign share or any interaction terms or controls beyond year fixed effects. A 1% reduction in trade costs raises exports to the US by 10%, an elasticity comparable to those reported by [McCaig, Pavcnik, and Wong \(2025\)](#) and within the range found in the literature ([Caliendo and Parro 2015](#)).

In Column (2) we include the foreign share and its interaction with BTA exposure, i.e., the export platform effect, and find that both enter with positive and significant coefficients. The positive direct effect of the foreign share is not surprising given that multinationals are typically more export-oriented than their domestic counterparts. To gauge the economic significance of the export platform effect, compare a sector at the (export-weighted) mean of the foreign share distribution to one a standard deviation above the mean. A 1% trade cost reduction raises exports by 10% at the mean, but by 15% one standard deviation above.

The 2SLS result in Columns (3) confirms the OLS findings qualitatively and points to a larger magnitude. More importantly, it gives the export platform effect a causal interpretation. The

first-stage F -statistic of 40 is well above conventional thresholds for weak instrument problems. In Column (4) we add two controls – the change in Chinese exports to the US and changes in rest-of-world tariff rates on Vietnamese exports, capturing competition and other trade policy shocks including Vietnam’s WTO accession in 2007. These controls slightly shrink our point estimate, but it remains highly significant and implies that a 1% trade cost reduction causes a 10% rise in exports for the average sector, but a 19% increase for a high foreign share sector.

In Column (5) we perform a natural placebo test and look for effects of the BTA prior to the signing of the agreement. Reassuringly, neither the BTA itself nor the export platform effect had a significant effect on export growth in this period, suggesting that the results in Columns (1) – (4) do not reflect pre-existing trends. Finally in Columns (6) and (7) we investigate the dynamics of this shock by splitting the post-BTA period into short run (2002-2006) and long run (2007-2015). Both the direct effect of the BTA and the export platform effect are larger in the short run, but even in the long run both coefficients are large and statistically significant. We therefore focus on the effects of the BTA over the entire period 2002-2015, and we will use the specification in Column (4), instrumenting for the foreign share and its interaction with the BTA and controlling for Chinese exports and rest-of-world tariffs, throughout the rest of this section.

Composition and spillovers

The positive export platform effects in Table 1 show that sectors with a greater foreign presence benefited disproportionately from the BTA. As the conceptual framework makes clear, this pattern

Table 1: Effects of Trade Liberalization on Exports

	Full Sample (2002-2015)				Pre BTA	Post BTA	
	OLS (1)	OLS (2)	IV (3)	IV (4)	1996-99 (5)	2002-06 (6)	2007-15 (7)
BTA	9.98*** (2.13)	9.90*** (1.59)	10.05*** (1.46)	9.87*** (1.47)	-0.77 (1.22)	10.40*** (1.10)	9.45*** (1.73)
Foreign Share		3.23*** (1.04)	5.72*** (0.55)	5.65*** (0.61)	0.69 (0.81)	4.61*** (0.57)	6.19*** (0.86)
Foreign Share \times BTA		16.76** (7.77)	34.97*** (7.48)	31.77*** (9.63)	-10.54* (6.34)	38.56*** (6.92)	27.23** (11.88)
Controls	\times	\times	\times	\checkmark	\checkmark	\checkmark	\checkmark
Observations	1286	1261	1261	1261	291	449	812
Industries	93	91	91	91	86	91	91
R^2	0.47	0.61	0.50	0.54	0.11	0.49	0.53
F -statistic	-	-	39.66	19.80	27.46	21.04	16.13

Notes: Each column reports coefficients from regressions of changes in log exports relative to 2000 on BTA exposure, foreign share, and their interaction. All specifications include year fixed effects and are weighted by industry exports in 2000. Columns (4)–(7) additionally include log changes in Chinese exports to the United States and changes in rest-of-world tariff rates on Vietnamese exports. Robust standard errors clustered by industry are reported in parentheses. F -statistics report the Kleibergen-Paap Wald F -statistic across the two endogenous variables.

could reflect the *composition* channel, whereby foreign firms are inherently more responsive to new export opportunities, or the *spillover* channel, whereby foreign presence raises the responsiveness of domestic firms, or both.

To distinguish between these channels, we turn to firm-level data from the VES, split the sample into domestic versus foreign-owned firms, and then aggregate to the sector level. In Table 2 we report estimates of (6) for each subsample using the same instruments, controls, and weights as in Table 1. The outcome variables we consider are log changes in the number of exporters, total employment, and total revenue.³ If only the composition channel is operating, we would expect larger direct effects for foreign firms but no interaction effects within either group. If spillovers are also present, we should see significant interaction effects for domestic firms as well.

The first row of Table 2 shows that the direct effects of the BTA on exporters, employment, and revenue are positive for both types of firm. For domestic firms, in Columns (1)-(3), these effects are smaller and less precisely measured. Consistent with the idea that foreign firms respond more strongly to new export opportunities, the results for foreign firms in Columns (4)-(6) are uniformly larger and significant at the 5% level. These estimates provide support for the composition channel.

The third row of Table 2 shows estimates of export platform effect within each group of firms: they show that the spillover channel also plays an important role. For domestic firms in Columns (1)-(3), we estimate large and significant export platform effects across the different outcomes. The

Table 2: Effects of Trade Liberalization by Firm Type

	Private Firms			Foreign Firms		
	Exporters (1)	Employment (2)	Revenue (3)	Exporters (4)	Employment (5)	Revenue (6)
BTA	0.27* (0.16)	1.01 (0.83)	1.42* (0.80)	2.53*** (0.36)	2.97*** (0.50)	1.93*** (0.68)
Foreign Share	-0.20 (0.25)	1.48*** (0.53)	0.58** (0.28)	0.09 (0.25)	-0.42 (0.31)	-1.31*** (0.43)
Foreign Share \times BTA	3.48** (1.77)	10.38** (4.70)	6.23* (3.61)	4.71** (2.08)	7.72*** (3.00)	3.31 (3.28)
Controls	✓	✓	✓	✓	✓	✓
Observations	592	1197	1197	623	1110	1095
Industries	68	86	86	72	81	80
R^2	0.48	0.26	0.59	0.43	0.39	0.56
F -statistic	8.53	20.83	20.83	25.43	22.72	24.79

Notes: Each column reports IV estimates, instrumenting for both F_i and $F_i T_i$ using Z_i and $Z_i T_i$. Dependent variables are the log change in the number of exporters, total employment, and total revenue within each ownership group at the four-digit industry level. All regressions include year fixed effects, controls for Chinese exports to the US and rest-of-world tariffs on Vietnam's exports, and are weighted by industry exports in 2000. Robust standard errors are clustered at the industry level.

³Employment and revenue data are available for all years 2000–2015, while data on the number of exporters are available for years 2000, 2002–2004. We focus on the number of exporters rather than export sales by firm type as the VES data on export sales contains substantial measurement error.

estimates in Column (3), for example, imply that a domestic firm at the mean foreign share that experiences a 1% cut in trade costs would increase its total revenues by 1.4%. The same cut in trade costs in a high foreign share sector would lead to a 3.3% increase in total revenues. The coefficients for foreign firms in Columns (4)-(6) are also positive, but in general somewhat smaller than those in (1)-(3), and in the case of revenue statistically insignificant. These results are consistent with the idea that the spillover channel primarily affects domestic firms.

Evidence of technology diffusion

The results in Table 2 point to positive spillovers from foreign to domestic firms. What generates these spillovers? A large literature has emphasized the possibility that technology diffuses from foreign to domestic firms through supply chain relationships, labor mobility, and direct observation of production techniques (Javorcik 2004; Buera and Oberfield 2020; Alfaro-Urena, Manelici, and Vasquez 2022). We now use the Technology and Competitiveness Survey to investigate this channel. We structure this analysis around the model of Buera and Oberfield (2020), in which technology diffusion involves three steps – exposure to new ideas, adaptation of these ideas through some innovative effort, and the resulting process or product modifications – a sequence that will also underpin the quantitative model in Section 2.

Table 3: Evidence of Technology Diffusion

	# MNC Connected (1)	Tech Transfer (2)	R&D (3)	Modify Technology (4)	# Successes (5)
BTA	1.52** (0.75)	0.08 (0.08)	0.11 (0.15)	0.13*** (0.03)	0.52*** (0.13)
Foreign Share	0.91*** (0.35)	0.13*** (0.03)	0.51*** (0.13)	0.02 (0.02)	0.03 (0.08)
Foreign Share × BTA	12.64*** (4.62)	1.93*** (0.43)	3.35*** (1.22)	1.00*** (0.21)	2.54*** (0.66)
Controls	✓	✓	✓	✓	✓
Observations	320	320	396	396	396
Industries	82	82	82	82	82
R ²	0.20	0.33	0.10	0.19	0.34
F-statistic	27.31	27.31	27.88	27.88	27.88

Notes: This table reports IV estimates of the effect of foreign firm presence on technology upgrading outcomes among private domestic Vietnamese manufacturing firms. Data are from the Vietnam Technology and Competitiveness Survey (TCS), 2010–2014, aggregated to the industry-year level. Column (1) “# MNC Connected” counts the mean number of foreign enterprises among a firm’s suppliers and customers. Column (2) “Tech Transfer” is the share of firms reporting any technology transfer from either domestic or foreign sources. Column (3) “R&D” share of firms conducting research and development activities. Column (4) “Modify Technology” indicates whether firms adapt existing technologies to their production processes; Column (5) “# Successes” measures the mean number of successful technology modifications per firm. All regressions include year fixed effects and controls for Chinese exports to the US and rest-of-world tariffs. Observations are weighted by industry export value in 2000. Robust standard errors clustered at the industry level in parentheses.

Table 3 reports estimates of (6) for domestic firms, using the same instruments, controls, and weights as above. Because the survey covers the period 2010–2014, outcomes are measured in levels rather than in changes. We find positive and statistically significant export platform effects at each stage. Domestic firms in high foreign share sectors report more supplier and customer relationships with foreign firms (Column 1), more technology transfers (Column 2), and more R&D activity (Column 3). These inputs translate into outputs: Columns (4) and (5) show that domestic firms in export platform sectors are more likely to modify their production processes, and complete more modification projects. In Appendix Table B1 we replicate this analysis for foreign-owned firms and find mixed results, consistent with the idea that foreign firms are closer to the technological frontier and have less to learn from their peers.

1.6 Taking Stock

We have studied the interaction of increases in export market access, in the form of the BTA, with differences in foreign firm presence, as measured by the foreign share in pre-BTA total revenues: we term this interaction the export platform effect.

Three findings emerge. First, we find that export platform effects are positive and quantitatively important: a 1% reduction in trade costs raises exports by 10% in a sector at the mean foreign share, but by 19% in a sector one standard deviation above. Second, we find that this positive interaction is the product of two forces: a composition channel driven by larger direct effects of increased export market access on foreign firms, and a spillover channel whereby domestic firms in export platform sectors also benefit. Third, evidence from the Technology and Competitiveness Survey points to a role for technology diffusion in generating these positive spillovers.

Our empirical results have two important shortcomings. First, the spillover channel we have explored combines both true externalities, like technology diffusion, and standard equilibrium interactions between firms, like competition in input and output markets. Disentangling these is challenging empirically but crucial from a policy perspective: only externalities provide a rationale for policies that target multinationals. Second, we can only identify relative effects across sectors. The central question of how the composition and spillovers channels we have identified ultimately shape Vietnam’s aggregate gains from the BTA, and its growth more broadly, remains open. The quantitative model we develop below addresses both these limitations.

2 A Theory of Multinational Production and Technology Diffusion

We now build a general equilibrium model to quantify the gains from trade and openness to multinationals. How much did Vietnam gain from the BTA, and how would these gains have been different had Vietnam been closed to foreign firms? To answer these questions, we embed the model of multinational production from [Ramondo and Rodríguez-Clare \(2013\)](#) within the multisector input-output structure of [Caliendo and Parro \(2015\)](#), and introduce technology diffusion building

on Buera and Oberfield (2020) and Cai and Xiang (2025). The resulting model implies a special role for export platforms in driving growth.

2.1 Environment

Time is continuous and denoted by t . The world economy is composed of countries $i = 1, \dots, N$ and sectors $s = 1, \dots, S$. We use subscripts to index countries and superscripts to index sectors, and where possible omit t . We follow the convention that j refers to a country where goods are produced, i refers to a destination country, and h refers to a country where a firm is headquartered. We refer to sector-origin-headquarter combinations as an (s, j, h) cell.

Each country is endowed with a fixed mass of workers l_j , who move freely between different sectors and supply labor in exchange for a common wage w_j . Within each sector, there is a unit continuum of varieties indexed by ω , and within each (s, j, h) cell there are a large number of perfectly competitive potential producers of each variety. At each point in time the distribution of productivity across these potential producers, which we term the technology frontier, is given. Over time, however, it will evolve endogenously as technologies diffuse across producers.

2.2 Production and trade

We model a roundabout production structure: in each sector, different varieties ω of intermediate goods are produced by heterogeneous firms using labor and composite intermediate goods from all sectors. Composite intermediate goods are in turn produced by aggregating the varieties ω within each sector.

Intermediate goods

An intermediate goods producer in sector s , headquartered in country h , can set up production in any country j , and may choose to operate in multiple countries. If such a firm hires l workers and purchases m^r units of composite intermediates from each sector r , they produce output according to

$$y_{jh}^s(\omega) = \left(\frac{z_{jh}^s(\omega)}{\delta_{jh}^s} \right) f_j^s(l, m^1, \dots, m^S), \quad (8)$$

where the production function f_j^s is common to all firms operating in sector s in country j regardless of their variety ω or headquarters location h . $z_{jh}^s(\omega)$ is the productivity of a firm headquartered in h and producing in j , which we model as a draw from a distribution. We refer to δ_{jh}^s as the multinational production cost, and normalize $\delta_{hh}^s = 1$.

We assume the production function f_j^s is Cobb-Douglas with a labor share γ_j^s and intermediate input shares $(1 - \gamma_j^s)\beta_j^{sr}$, with the sum of β_j^{sr} equal to one. We denote the cost of a unit bundle of

inputs for this production function by x_j^s :

$$x_j^s = w_j^{\gamma_j^s} \left(\prod_{r=1}^S (P_j^r)^{\beta_j^{sr}} \right)^{1-\gamma_j^s}, \quad (9)$$

where P_j^r is the price of composite intermediates from sector r in country j .

Composite intermediate goods

A perfectly competitive representative firm in each sector s and country i aggregates varieties ω into a composite intermediate good, purchasing each variety from its lowest cost supplier. As in [Caliendo and Parro \(2015\)](#), this aggregator has a constant elasticity of substitution (CES) form with elasticity σ . Composite intermediates may be used as inputs or consumed by households.

International trade

All firms operating in country j can ship goods to country i subject to iceberg trade costs τ_{ijh}^s , which may depend on the headquarters location h . A firm in (s, j, h) producing variety ω that ships one unit to country i incurs costs,

$$c_{ijh}^s(\omega) = \frac{\kappa_{ijh}^s}{z_{jh}(\omega)}, \quad \kappa_{ijh}^s \equiv x_j^s \delta_{jh}^s \tau_{ijh}^s. \quad (10)$$

Given our assumption of perfect competition, such a firm will only ship to destination i if it is the lowest cost supplier of variety ω in that country, and if so it will charge a price exactly equal to $c_{ijh}^s(\omega)$.

Let Ω_{ijh}^s be the subset of varieties for which the lowest cost supplier to i is a firm in (s, j, h) . Then the CES production technology for composite intermediates implies trade shares π_{ijh}^s and price indices P_i^s ,

$$\pi_{ijh}^s = \int_{\Omega_{ijh}^s} \left(\frac{c_{ijh}^s(\omega)}{P_i^s} \right)^{1-\sigma} d\omega, \quad P_i^s = \left(\sum_j \sum_h \int_{\Omega_{ijh}^s} c_{ijh}^s(\omega)^{1-\sigma} d\omega \right)^{\frac{1}{1-\sigma}}. \quad (11)$$

We also define i 's expenditure on goods produced in j as a share of its expenditure on goods produced by firms headquartered in h :

$$\psi_{ijh}^s = \frac{\pi_{ijh}^s}{\sum_{j'} \pi_{ij'h}^s}. \quad (12)$$

2.3 Technology diffusion

Each firm headquartered in h producing variety ω is characterized by a vector of potential productivities $z_h(\omega) = [z_{jh}(\omega)]_j$, drawn from a distribution G_h^s that we call the *technological frontier*. New

ideas—vectors of potential productivities across locations—arrive according to a Poisson process. Each new idea z can be decomposed as $z = ab^\rho$, where a is a vector-valued *original component* drawn from an exogenous distribution, and b is a scalar *learning component* drawn from an endogenous source distribution $H_h^s(b)$. The parameter $\rho \in [0, 1)$ is the *learning intensity*: it captures the extent to which existing production technologies inform new ideas.

Following [Cai and Xiang \(2025\)](#), we specify the distribution of a as follows.

Assumption 1. *For sector s in country h , the arrival rate of ideas with an original component greater than a in at least one dimension is*

$$A_h^s(a) = \alpha_h^s \left(\sum_j a_j^{\frac{-\theta}{1-\eta}} \right)^{1-\eta}, \quad (13)$$

for any $a \gg 0$. Here $\alpha_h^s > 0$ is a time-varying, exogenous arrival intensity shifter, $\theta > 0$ determines the shape of the right tail of the distribution, and $\eta \in [0, 1)$ determines the correlation between different elements of the vector of productivity draws.

Taking the distribution of b as given, [Assumption 1](#) implies that the rate of arrival of ideas which are weakly better than z in at least one dimension is,

$$\tilde{A}_h^s(z) = \alpha_h^s \left(\sum_j z_j^{\frac{-\theta}{1-\eta}} \right)^{1-\eta} \int_0^\infty x^{\rho\theta} dH_h^s(x). \quad (14)$$

The final integral captures the role of learning in generating improvements over the vector z .

The source distribution H_h^s describes how potential producers learn from existing ones. Following the “learning from workers” specification of [Buera and Oberfield \(2020\)](#), we assume that potential producers in sector s headquartered in h sample from all production technologies in use in country h —that is, from technologies employed by all firms producing locally, regardless of where those firms are headquartered. We use k to index the headquarters of these existing producers, reserving h for the headquarters of the learning firm. The resulting source distribution is,

$$H_h^s(b) = \sum_i \sum_k \int_{\Omega_{ihk}^s} \ell_{ihk}^s(\omega) \mathbb{1}\{z_{hk}^s(\omega) \leq b\} d\omega, \quad (15)$$

where $\ell_{ihk}^s(\omega)$ is the share of total employment in sector s in country h that is devoted to the production of variety ω by firms headquartered in k and selling to country i . Sampling probability is proportional to employment, so larger firms contribute more to the pool of ideas.

The technological frontier

Before stating the main result, define r_{ihk}^s to be the share of total revenues in sector s in country h earned by firms headquartered in k , producing in h , and selling to destination i . Combining [\(14\)](#) and [\(15\)](#) then yields [Theorem 1](#).

Theorem 1. *At each time t , the technological frontier G_h^s is a multivariate Fréchet distribution with shape parameter θ and correlation parameter η . We denote the scale parameter in sector s in country h by λ_h^s , and refer to this as the level of technology. Its law of motion is,*

$$\dot{\lambda}_h^s = \tilde{\alpha}_h^s \sum_k \sum_i r_{ihk}^s \left(\frac{\lambda_k^s}{\pi_{ihk}^s (\psi_{ihk}^s)^{-\eta}} \right)^\rho, \quad (16)$$

where

$$\tilde{\alpha}_h^s = \alpha_h^s \left(\frac{\Gamma(1 - \rho - (\frac{\sigma-1}{\theta}))}{\Gamma(1 - (\frac{\sigma-1}{\theta}))} \right) \quad (17)$$

is an exogenous shifter and Γ is the Gamma function.

Theorem 1 contains two results. The first is that in each country and sector, productivity distributions are multivariate Fréchet, exactly as in [Ramondo and Rodríguez-Clare \(2013\)](#). But whereas in that paper these distributions are exogenous, here they are the endogenous result of the process of technology diffusion described above. This result is key to the model's tractability. The second is that the scale parameters of these distributions—the levels of technology—evolve over time according to (16), in a way that depends on trade flows and revenue shares. This creates the possibility of spillovers from multinational firms operating in a given country to their domestic peers, a mechanism we discuss in [Section 2.5](#).

2.4 Equilibrium

Theorem 1 implies natural expressions for trade shares and price indices,

$$\psi_{ijh}^s = \left(\frac{\kappa_{ijh}^s}{K_{ijh}^s} \right)^{-\frac{\theta}{1-\eta}}, \quad K_{ih}^s \equiv \left(\sum_j (\kappa_{ijh}^s)^{\frac{-\theta}{1-\eta}} \right)^{\frac{1-\eta}{-\theta}} \quad (18)$$

$$\pi_{ijh}^s = \psi_{ijh}^s \left(\frac{\lambda_h^s (K_{ih}^s)^{-\theta}}{\sum_{h'} \lambda_{h'}^s (K_{ih'}^s)^{-\theta}} \right) \quad (19)$$

$$P_i^s = \Gamma \left(1 - \left(\frac{\sigma-1}{\theta} \right) \right)^{\frac{1}{1-\sigma}} \left(\sum_h \lambda_h^s (K_{ih}^s)^{-\theta} \right)^{\frac{-1}{\theta}}. \quad (20)$$

Given a set of technology levels λ , define a *static equilibrium* to be a collection of trade flows, price indices, revenues, expenditures and wages satisfying (18) – (20). Revenues and expenditures are defined by,

$$R_j^s = \sum_h \sum_i \pi_{ijh}^s X_i^s, \quad X_j^s = \alpha_j^s w_j l_j + \sum_r (1 - \gamma_j^r) \beta_j^{rs} R_j^r \quad (21)$$

where α_j^s is the weight of sector s in final consumption in country j . Finally the labor market must clear, implying,

$$w_j l_j = \sum_s \gamma_j^s R_j^s. \quad (22)$$

Let us use \mathcal{E} to denote a static equilibrium, and let $\mathcal{E}(t)$ denote the static equilibrium at time t . Then, given an initial set of technology levels $\lambda(0)$, a *dynamic equilibrium* is a path for technology levels and static equilibria $\{\lambda(t), \mathcal{E}(t)\}_{t>0}$, such that at every point in time $\mathcal{E}(t)$ is a static equilibrium given $\lambda(t)$, and $\lambda(t)$ solves (16) given the initial condition $\lambda(0)$.

Balanced growth

To generate balanced growth, we assume that for every country and sector α_h^s grows at a constant rate $g_\alpha > 0$. We show in Appendix A that all technology parameters λ_h^s then grow at rate $g_\lambda = (1 - \rho)^{-1}g_\alpha$, and that along the balanced growth path these solve,

$$\tilde{\lambda}_h^s = \left(\frac{\tilde{\alpha}_h^s(0)}{g_\lambda} \right) \sum_k \sum_i r_{ihk}^s \left(\frac{\tilde{\lambda}_k^s}{\pi_{ihk}^s (\psi_{ihk}^s)^{-\eta}} \right)^\rho, \quad \tilde{\lambda}_h^s \equiv \lambda_h^s \exp(-g_\lambda t). \quad (23)$$

2.5 Discussion

The model contains two forces that can generate the empirical patterns documented in Section 1: multinational mobility, governed by η , and learning from export platforms, governed by ρ . Throughout, we focus on a single country j —the model counterpart of Vietnam—and abstract from general equilibrium effects by holding all wages and price indices constant.

Multinational mobility

To isolate the role of *multinational mobility*, set $\rho = 0$ so that technologies are exogenous. Consider a stylized environment in which: (i) firms headquartered in j face prohibitively high costs of operating outside j ; and (ii) for firms headquartered in other countries $h \neq j$, affiliates producing in j represent a negligible share of their total production. These two assumptions create a clean distinction between domestic versus foreign firms that aligns with the empirical setting of Section 1.

Mimicking the BTA reform studied in Section 1, suppose firms in sector s operating in j experience a small reduction in the iceberg trade cost of selling to market i . Denote this change by $T^s \equiv -\Delta \log \tau_{ijh}^s$. The implied (partial equilibrium) change in exports for domestic firms is

$$\Delta \log \pi_{ijj}^s = \theta T^s.$$

For foreign firms, headquartered in any $h \neq j$, the change in exports is instead

$$\Delta \log \pi_{ijh}^s = \left(\frac{\theta}{1 - \eta} \right) T^s.$$

When $\eta > 0$, multinational exports are more sensitive to trade-cost shocks (and, more generally, to any shock affecting relative profitability across locations). Intuitively, multinationals respond along

an additional margin—choosing where to produce—and the strength of this margin depends on the degree of correlation in their productivity draws across locations, η .

Learning from export platforms

Moving beyond the static role of multinationals, we now isolate the role of learning. Set $\eta = 0$ and focus on the technology of firms headquartered in j in sector s , $\tilde{\lambda}_j^s$. In the long run (i.e., along the balanced growth path), $\tilde{\lambda}_j^s$ is pinned down by (23), which can be written as

$$\tilde{\lambda}_j^s = \left(\frac{\tilde{a}_j^s(0)}{g_\lambda} \right) \times \left(\frac{\tilde{\lambda}_j^s}{\pi_{jjj}^s} \right)^\rho \times \left(\sum_i \sum_h r_{ijh}^s \left(\delta_{jh}^s \tau_{ijh}^s \left(\frac{P_i^s}{P_j^s} \right) \right)^{\rho\theta} \right). \quad (24)$$

The first term reflects j 's exogenous ability to generate new ideas in sector s . The second term captures domestic learning from firms headquartered in, operating in, and selling to j . The third term is our focus: it captures what j -headquartered firms learn from exposure to export markets i and from multinationals headquartered in foreign countries h , weighting each destination-headquarter pair (i, h) by its revenue share r_{ijh}^s .

Equation (24) implies that “tougher” destination-headquarter pairs generate larger spillovers, where toughness depends on multinational production costs δ , iceberg trade costs τ , and the intensity of competition, as captured by relative price indices. The intuition is selection: firms that can profitably serve tough markets must be highly productive and therefore embody high-quality ideas. Because multinationals that export must bear both multinational production costs and trade costs, they are more intensely selected—and thus more productive—than either multinationals whose production in j is sold only to j , or domestic firms that export.

Export platforms are therefore an especially powerful channel for technology diffusion. Any shock that expands the role of export-platform activity in j (i.e., raises the revenue share of exporting multinationals operating locally) in the long run results in higher productivity among j -headquartered firms. The strength of this mechanism is governed by ρ . When $\rho = 0$, learning is absent and productivities are exogenous, while as $\rho \rightarrow 1$ new ideas are increasingly drawn from the toughest markets and export platforms play an increasingly important role.

Summary

The two mechanisms, multinational mobility and learning from export platforms, both imply that a shock to trade costs such as the BTA will have different effects in sectors with a strong foreign firm presence. But they have distinct empirical implications. Multinational mobility can generate a strong aggregate export response in foreign-intensive sectors, but it cannot generate positive spillovers to domestic firms. Learning from export platforms can generate precisely these spillovers by raising domestic productivity as export-platform activity expands. Our calibration strategy will discipline η and ρ by matching these patterns jointly.

3 Calibration

We now calibrate the model. As detailed below, our calibration is stylized rather than comprehensive, and we ignore aspects of the data that are not central to the mechanisms we focus on.

We consider a world economy composed of four countries: Vietnam (VN), the United States (US), a composite East Asia (EA) composed of Japan, Korea, Singapore and Taiwan—i.e., the leading sources of multinational firms in Vietnam—and a composite Rest of World (RoW). We take our sector classification from the OECD’s IO tables and compress all nontradables into a single sector, leaving us with 24 sectors in total.

3.1 Hat algebra

Our model has a large number of unknown fundamentals: iceberg trade costs, multinational production costs, and so on. We deal with these using the ‘hat algebra’ approach that is standard in the quantitative trade literature (Dekle, Eaton, and Kortum 2007). Appendix Theorem 1 spells out our procedure in full. Briefly, we show how to solve for the effects of counterfactual changes to model fundamentals, using only data from a baseline equilibrium and values for the parameters (θ, η, ρ) , without explicitly recovering the baseline values of those fundamentals.

We take the baseline equilibrium to be the world economy in the year 2000, just before the BTA reform that we studied in Section 1. Most of the data required by our hat algebra procedure is standard and can be read directly from the OECD’s IO tables, but two points are worth describing in detail.

Baseline trade flows

The OECD reports trade flows between producers and buyers, rather than the full matrix of flows between headquarters, producers, and buyers required by our model. First, we abstract away from multinational production by any countries other than East Asia, and assume that for East Asian firms the only feasible locations for production are Vietnam or East Asia itself. This immediately reduces the missing data problem to the task of filling in π_{ijh}^s for $j = VN$ and $h \in \{VN, EA\}$. Second, we simplify the problem of inferring these trade shares by parameterizing bilateral trade costs as $\tau_{ijh}^s = \tau_h^s \tau_{ij}^s$ for $i \neq j$. This assumption allows for a shifter τ_h^s in the trade costs facing East Asian firms producing in Vietnam, but restricts this shifter not to vary across destination markets i . Third, we note that from the VES data introduced in Section 1, we can observe the revenues of multinationals and domestic firms in each sector s , separated into sales made in Vietnam versus abroad. Combined with our assumption on τ_{ijh}^s , this data pins down the full set of trilateral trade shares π_{ijh}^s in each sector.

Baseline technology diffusion

Solving the technology diffusion block of the model requires information on idea flows between firms with different headquarters and serving different markets in the baseline equilibrium. These flows can be inferred from the structure of the model alongside observations on the technology levels λ_h^s in each country h and sector s in the baseline equilibrium; but, of course, these technology levels are not directly observable.

We therefore use information on trilateral flows and the ratio of unit costs to price indices for each country and sector to back out λ_h^s . The model implies that

$$\lambda_h^s \propto \left(\frac{x_h^s}{P_h^s} \right)^\theta \pi_{hhh}^s (\psi_{hhh}^s)^\eta, \quad (25)$$

where x_h^s represents the unit costs defined in (9) and P_h^s is the price index defined in (20). Intuitively, if a country h is able to capture a large share of its own market in sector s despite relatively high unit costs, it must be that its firms are highly productive, as captured by their technology level λ_h^s .

To operationalize (25), we read P_h^s from the World Bank's International Comparison Program (World Bank 2008) and construct unit costs x_h^s using IO coefficients from the OECD and real wages w_h . In turn we obtain real wages by equating them with real value added per capita, which we then infer from the OECD's IO tables.

3.2 Internally calibrated parameters

Three parameters remain to be calibrated: the productivity dispersion parameter θ , the correlation parameter η , and the learning intensity parameter ρ . Our calibration strategy builds on the discussion in Section 2.5, which highlighted the role of multinational mobility and learning from export platforms. We feed the model the actual BTA tariff changes, simulate regressions that parallel our empirical specifications in Section 1, and choose (θ, η, ρ) to match the estimated coefficients. We target the direct effects of the BTA and its interaction with the foreign share on two outcomes—Vietnam's exports to the US and domestic firm revenues. This gives us an overidentified system with three parameters targeting four moments.

Identification

The parameter θ disciplines the direct effect of the BTA on exports and on the total revenues of domestic firms. When θ is large, productivity dispersion across firms is low and trade flows—and thus revenues—are highly sensitive to changes in trade costs.

In line with the discussion in Section 2.5, the parameters η and ρ both play a role in the interaction between the BTA and the foreign share. But the mechanisms through which they operate are very different: η determines how responsive multinationals are to the BTA, while ρ determines how far these changes in multinational presence create spillovers to domestic firms.

Thus, targeting both the interaction effect for total exports and the interaction effect for domestic firm revenues allows us to separately identify η and ρ .

Implementation

We interpret the year 2000 equilibrium as the pre-BTA balanced growth path. We aggregate the BTA-induced tariff changes from Section 1 to the level of our model’s 23 tradable sectors and feed them to the model as permanent reductions in the iceberg cost τ_{ij}^s of shipping from Vietnam to the United States.

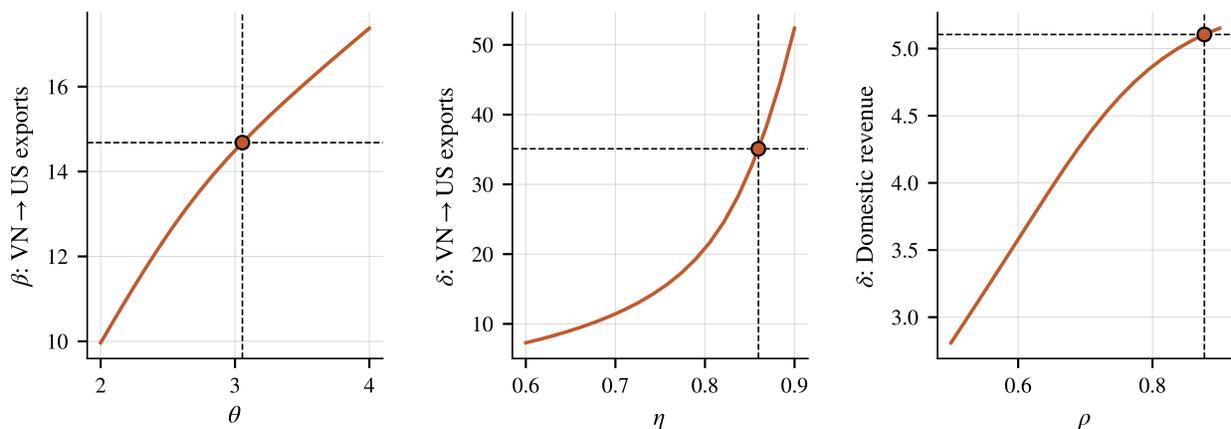
We solve the model for the full dynamic transition over $T = 15$ periods following the shock. Within each period, we solve for a static equilibrium—wages, prices, trade flows—and then update technology levels λ according to the law of motion (16). After solving, we compute log changes in total exports and domestic firm revenues by sector between $t = 0$ and $t = 15$. We then run the same weighted least squares regression as in the empirics: the outcome regressed on the tariff change T_s , the pre-BTA foreign share F_s , and their interaction $T_s \times F_s$, weighted by baseline sector size. We choose (θ, η, ρ) to minimize the sum of squared (percentage) deviations between model-implied and empirical regression coefficients.

Results

Table 4 reports the calibrated parameter values and the model’s fit to targeted moments, as well as to regression coefficients from Section 1 that were not targeted in our calibration. Figure 3 illustrates our identification arguments by showing how three of the target moments vary with the three internally calibrated parameters. Our calibration yields $\theta = 3.06$, $\eta = 0.86$, and $\rho = 0.88$.

Three features of these estimates stand out. First, with a nonzero correlation parameter η , the

Figure 3: Model moments and parameters



Notes: Each panel shows one how model-implied regression coefficient varies with one parameter while holding the other fixed. Coefficients labelled β represent the direct effects of the BTA. Coefficients labelled δ represent BTA-foreign share interaction effects.

productivity dispersion parameter θ is no longer equivalent to the trade elasticity as in [Eaton and Kortum \(2002\)](#). In fact, [Ramondo and Rodríguez-Clare \(2013\)](#) make the point that θ will generally be below the trade elasticity; since typical estimates of this elasticity range from 4 to 8, our estimate of $\theta = 3.06$ is consistent with this. Second, our estimate of $\eta = 0.86$ is higher than the value of 0.55 reported by [Arkolakis et al. \(2018\)](#), implying that we find multinational firms are relatively mobile: ultimately this finding is driven by the large interaction coefficients we found in the data. Third, we also find a large value for the learning intensity $\rho = 0.88$: [Buera and Oberfield \(2020\)](#) calibrate a value of 0.60, [Cai, Caliendo, et al. \(2025\)](#) provide a similar model of learning with a combined learning intensity of 0.62, and [Cai and Xiang \(2025\)](#) estimate 0.73. This is, to our knowledge, the first estimate of the learning intensity (in the context of international trade) that leverages plausibly

Table 4: Model Calibration: Parameters and Moments

<i>Panel A: Calibrated parameters</i>		
Parameter		Value
θ : Productivity dispersion		3.055
η : Multinational productivity correlation		0.8597
ρ : Learning intensity		0.8772
<i>Panel B: Targeted moments</i>		
Moment	Data	Model
<i>Direct effects (β)</i>		
VN→US exports	9.872 (1.469)	14.681
Domestic revenue	1.422 (0.798)	1.321
<i>Interaction effects (δ)</i>		
VN→US exports	31.768 (9.630)	35.098
Domestic revenue	6.226 (3.612)	5.105
<i>Panel C: Untargeted moments</i>		
Moment	Data	Model
<i>Direct effects (β)</i>		
Domestic employment	1.009 (0.833)	1.321
Domestic exporters	0.271 (0.162)	2.369
Foreign revenue	1.929 (0.684)	7.245
Foreign employment	2.969 (0.499)	7.245
Foreign exporters	2.534 (0.358)	10.100
<i>Interaction effects (δ)</i>		
Domestic employment	10.380 (4.704)	5.105
Domestic exporters	3.478 (1.768)	23.178
Foreign revenue	3.314 (3.282)	2.251
Foreign employment	7.722 (2.995)	2.251
Foreign exporters	4.705 (2.077)	7.923

Notes: Panel A show the parameter values that minimize the loss function described in the text. Panel B shows target moments and C untargeted moments, with standard errors in parentheses. All target moments are from regressions of an outcome variable on BTA exposure, the foreign share, and their interaction.

exogenous shifts in the source distribution of ideas.

Turning to model fit, the model matches the targeted regression coefficients reasonably well but not perfectly: the model generally overshoots the effects of the BTA on total exports but undershoots on domestic firm revenues. Among untargeted moments, the model generates qualitatively correct direct and interaction effects for most outcomes. The most noticeable failure is on the number of exporters, where the model-implied direct and interaction coefficients are consistently too large; but given that Eaton-Kortum-style trade models are not built to capture such extensive margin shifts, this shortcoming is not too surprising.

4 The Aggregate Consequences of Export Platforms

With the calibrated model in hand, we are now in a position to quantify the dynamic, general equilibrium effects of the BTA on Vietnam’s growth. We take as the baseline the world economy in 2000, and feed the model the changes in Vietnam’s iceberg costs of shipping to the United States caused by the BTA. We focus on the effect on Vietnamese welfare, defined below, and — surprisingly — find that while Vietnam enjoys large short-run gains from the BTA, in the long run these gains turn to losses.

Welfare

We solve the model for the full transition following the BTA and measure welfare as the present discounted value of per-period equivalent variation, defining the per-period welfare effect as the difference in real wage changes between the counterfactual (BTA) and baseline (no-BTA) paths,

$$\Delta W_t = \frac{\hat{w}_t^{\text{cf}}}{\prod_s (\hat{P}_t^{s,\text{cf}})^{\alpha^s}} - \frac{\hat{w}_t^{\text{base}}}{\prod_s (\hat{P}_t^{s,\text{base}})^{\alpha^s}},$$

where hats denote proportional changes at period t relative to the initial balanced growth path and we suppress the country subscript (all welfare results refer to Vietnam unless noted). The baseline path is not trivial: Vietnam’s technology levels grow even absent the BTA, so baseline wages and prices evolve over the transition. We then compute the discount-rate-weighted average $\overline{\Delta W}$ with a baseline annual discount rate of $r = 5\%$.⁴

4.1 Results

Vietnam gains from the BTA in present value terms – welfare rises by 1.24%. But these gains decay over time. Panel (a) of Figure 4 traces out this evolution over 15 years following the BTA. On impact ($t = 1$), technology levels do not respond and the BTA effect equals the static gain of 3.39%. But as technology levels adjust to the new trade environment, the welfare gain shrinks: to 2.37% by year 3,

⁴Specifically, $\overline{\Delta W} = \sum_{t=1}^T \left(\frac{1}{1+r}\right)^{t-1} \Delta W_t / \sum_{t=1}^T \left(\frac{1}{1+r}\right)^{t-1}$.

1.66% by year 5, and 0.23% by year 11, crossing zero between years 11 and 12 and reaching -1.22% by year 15.

We henceforth focus on the long-run effects in year 15, and investigate why they turn negative. To isolate the roles of multinational production and endogenous technology diffusion, we run three pairs of counterfactual experiments. Each pair compares a baseline path to a path with BTA tariff cuts, while varying whether multinationals are active and whether technology levels are allowed to respond. Scenario A is the full BTA with all mechanisms active. Scenario B imposes a prohibitive multinational production cost so that no foreign firms can produce in Vietnam, isolating the pure trade-cost channel. Scenario C holds technology levels fixed at baseline levels, giving the textbook static gains from trade. Within each pair, differencing the counterfactual from the baseline yields an additive decomposition of the total BTA effect into a contribution from multinationals and a dynamic contribution.

Multinationals

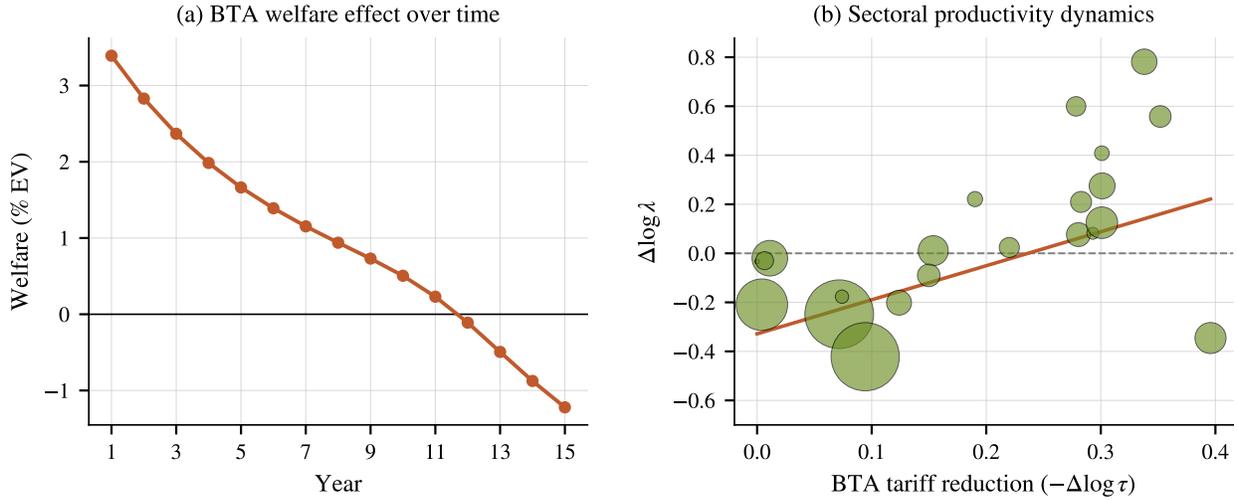
The BTA's effects on Vietnam operate almost entirely through multinationals. In Scenario B, in which multinational production is shut down, the BTA generates a welfare gain of just 0.22%. Our calibration implies multinationals are much more mobile than domestically owned producers: they therefore play an outsized role following the BTA shock in reallocating employment and exports across sectors to exploit changes in Vietnam's comparative advantage. With multinationals present, Vietnam's exports to the US rise by 2.25 log points (roughly 9.5-fold) as a result of the BTA, and 90% of this increase is attributable to multinationals. But greater multinational mobility *per se* does not explain the long-run losses from the BTA: these must ultimately be driven by endogenous shifts in productivity.

Technology diffusion

Decomposing the long-run effect into a static and dynamic component reveals that technology diffusion is the quantitatively dominant force. Holding technology levels fixed – Scenario C – the BTA generates straightforward gains of 3.39%. In line with the predictions of standard static trade models, tariff cuts expand export demand and raise real wages in Vietnam. But the endogenous response of productivity in Vietnam more than offsets these gains in the long run, contributing -4.61 percentage points. The aggregate welfare effect is thus almost entirely governed by how the BTA reshapes technology diffusion.

The explanation for these long-run losses ultimately comes from general equilibrium changes in wages. While multinationals are more responsive to the positive direct effects of the BTA, they also contract more in response to wage increases. On impact, the BTA raises nominal wages in Vietnam by 9.8%. Sectors that experience only modest tariff decreases thus face a negative shock on net, and the result is that multinational production contracts, setting off a gradual decline in Vietnamese technology levels.

Figure 4: Effects of the BTA



Notes: Panel (a) plots the per-period welfare effect of the BTA on Vietnam, measured as the difference in equivalent variation between the BTA and no-BTA paths. Panel (b) plots BTA tariff reductions ($-\Delta \log \tau$) against changes in Vietnam's knowledge stock relative to the no-BTA baseline ($\Delta \log \lambda$) for each tradable sector. Bubble size is proportional to domestic revenue weight. The line shows the weighted least squares fit, with weights proportional to domestic revenue.

Panel (b) of Figure 4 illustrates this pattern. Each bubble represents a tradable sector; the horizontal axis plots the BTA tariff reduction, the vertical axis plots the long-run change in Vietnam's productivity relative to the no-BTA baseline, and bubble size is proportional to domestic revenue. Sectors with large tariff cuts (other manufacturing, rubber and plastics, computer and electronics, fabricated metal products) cluster in the upper right: technology levels rise faster than they would absent the BTA. Sectors with small tariff cuts (agriculture, food products, basic metals) cluster in the lower left: the nominal wage increase dominates, multinational production contracts, and technology diffusion slows. The aggregate slowdown in technology growth is driven by the lower-left cluster. Of 23 tradable sectors, 12 see faster technology growth under the BTA (average 0.27, accounting for 27% of revenue weight) and 11 see slower growth (average -0.29 , accounting for 73%). Food products alone account for 16.3% of consumption weight, agriculture for 9.2%. The revenue-weighted average change in technology levels is -0.138 .

4.2 Taking stock

Three results emerge from the counterfactual analysis. First, Vietnam's gains from the BTA are positive and large in present value terms: welfare rises by 1.24%. Second, multinationals and technology diffusion are jointly the decisive factor – without them, the BTA generates a negligible 0.22% gain. Third, the long-run effects of technology diffusion are negative: general equilibrium wage increases cause multinational production to contract in sectors with small tariff cuts, and the resulting slowdown in technology growth more than offsets the static gains.

These results are preliminary and may evolve as we refine our empirical strategy and calibration. But the central lesson of this section – that the long-run, general equilibrium effects of a trade shock may look quite different from what partial equilibrium responses would suggest – is, we believe, a robust insight of the model.

5 Conclusion

We have studied how export platforms shaped Vietnam’s growth, and in particular its gains from the US-Vietnam Bilateral Trade Agreement, or BTA. Using the BTA as a source of exogenous variation in export market access together with a shift-share instrument for foreign presence, we have shown that sectors with greater pre-reform multinational activity respond more strongly to improved export-market access, and that these amplified responses extend to domestic firms operating in foreign-intensive sectors. We interpret these patterns as evidence that export platforms matter not only because multinationals are more mobile and export-oriented, but also because they facilitate technology diffusion to domestic producers. A quantitative general equilibrium model disciplined by our reduced-form estimates provides a framework for assessing how these mechanisms map into aggregate outcomes. We find that Vietnam’s gains from the BTA were sizable in present value terms, but declined over time due to the interaction of multinational mobility, technology diffusion, and general equilibrium changes in wages.

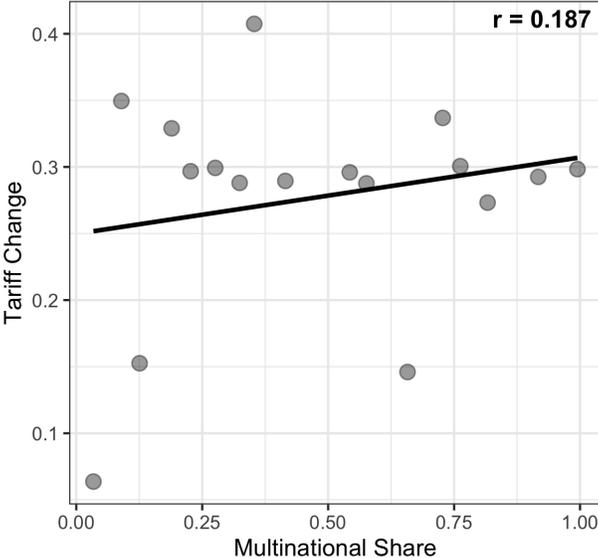
References

- Aitken, Brian J and Ann Harrison (1999). "Do domestic firms benefit from direct foreign investment? Evidence from Venezuela". In: *American Economic Review* 89.3, pp. 605–618.
- Alfaro-Urena, Alonso, Isabela Manelici, and Jose P. Vasquez (July 2022). "The Effects of Joining Multinational Supply Chains: New Evidence from Firm-to-Firm Linkages". In: *The Quarterly Journal of Economics* 137 (3), pp. 1495–1552. ISSN: 0033-5533.
- Arkolakis, Costas, Natalia Ramondo, Andrés Rodríguez-Clare, and Stephen Yeaple (2018). "Innovation and production in the global economy". In: *American Economic Review* 108.8, pp. 2128–2173.
- Blonigen, Bruce and Alyson Ma (2007). *Please pass the catch-up the relative performance of chinese and foreign firms in chinese exports*.
- Buera, Francisco J. and Ezra Oberfield (Jan. 2020). "The Global Diffusion of Ideas". In: *Econometrica* 88 (1), pp. 83–114.
- Cadestin, Charles, Koen De Backer, Isabelle Desnoyers-James, Sébastien Miroudot, Ming Ye, and Davide Rigo (2018). "Multinational enterprises and global value chains: New Insights on the trade-investment nexus". In: *OECD Science, Technology and Industry Working Papers* 2018.5, 0_1–36.
- Cai, Sheng, Lorenzo Caliendo, Fernando Parro, and Wei Xiang (2025). "Mechanics of Spatial Growth". In: *Working Paper*.
- Cai, Sheng and Wei Xiang (2025). "Multinational production, technology diffusion, and economic growth". In: *Journal of International Economics* 157, pp. 104–137.
- Caliendo, Lorenzo and Fernando Parro (2015). "Estimates of the Trade and Welfare Effects of NAFTA". In: *The Review of Economic Studies* 82.1, pp. 1–44.
- Chandra, Vandana and Sashi Kolavalli (2006). *Chandra V (ed) Technology, adaptation, and exports: How some countries got it right, chapter 1 in technology, adaptation and exports*.
- Dekle, Robert, Jonathan Eaton, and Samuel Kortum (2007). "Unbalanced trade". In: *American Economic Review* 97.2, pp. 351–355.
- Drozdoff, Natalia, Brian McCaig, and Nina Pavcnik (2025). "Foreign Manufacturing Jobs and Women: Impacts on Work, Marriage, and Childbirth". Working paper.
- Eaton, Jonathan and Samuel Kortum (2002). "Technology, geography, and trade". In: *Econometrica* 70.5, pp. 1741–1779.
- Fukase, Emiko (Jan. 2013). "Export Liberalization, Job Creation, and the Skill Premium: Evidence from the US–Vietnam Bilateral Trade Agreement (BTA)". In: *World Development* 41, pp. 317–337. ISSN: 0305-750X. DOI: [10.1016/j.worlddev.2012.05.035](https://doi.org/10.1016/j.worlddev.2012.05.035).
- Harrison, Ann and Andrés Rodríguez-Clare (2010). "Trade, foreign investment, and industrial policy for developing countries". In: *Handbook of development economics* 5, pp. 4039–4214.
- Hoang, Trang, Devashish Mitra, and Hoang Pham (2025). "The Effect of Export Market Access on Labor Market Power: Firm-level Evidence from Vietnam". Working paper.

- Javorcik, Beata Smarzynska (2004). “Does Foreign Direct Investment Increase the Productivity of Domestic Firms? In Search of Spillovers Through Backward Linkages”. In: *American Economic Review* 94.3, pp. 605–627.
- Keller, Wolfgang (2022). “Knowledge spillovers, trade, and foreign direct investment”. In: *Oxford Research Encyclopedia of Economics and Finance*.
- Keller, Wolfgang and Stephen R Yeaple (2009). “Multinational enterprises, international trade, and productivity growth: firm-level evidence from the United States”. In: *The review of economics and statistics* 91.4, pp. 821–831.
- McCaig, Brian (2011). “Exporting out of poverty: Provincial poverty in Vietnam and US market access”. In: *Journal of International Economics* 85.1, pp. 102–113.
- McCaig, Brian and Nina Pavcnik (2018). “Export markets and labor allocation in a low-income country”. In: *American Economic Review* 108.7, pp. 1899–1941.
- McCaig, Brian, Nina Pavcnik, and Woan Foong Wong (Dec. 2025). “Foreign and Domestic Firms: Long Run Employment Effects of Export Opportunities”. In: *The Review of Economics and Statistics*, pp. 1–45. ISSN: 0034-6535. DOI: [10.1162/REST.a.1683](https://doi.org/10.1162/REST.a.1683).
- Mitra, Devashish, Hoang Pham, and Beyza Ural Marchand (2026). “Exports and Intergenerational Mobility”. Working paper.
- Newman, Carol, John Rand, Finn Tarp, Neda Trifković, and Tue Anh Nguyen (2015). *Firm-Level Technology and Competitiveness in Vietnam*. Tech. rep. Central Institute for Economic Management (CIEM), General Statistics Office (GSO), Development Economics Research Group (DERG) University of Copenhagen (UoC).
- Ramondo, Natalia and Andrés Rodríguez-Clare (2013). “Trade, multinational production, and the gains from openness”. In: *Journal of Political Economy* 121.2, pp. 273–322.
- STAR-Vietnam (2003). *An Assessment of the Economic Impact of the United States–Vietnam Bilateral Trade Agreement*. Tech. rep. Accessed on 2025-10-24. Hanoi, Vietnam: Support for Trade Acceleration Project (STAR-Vietnam).
- Tintelnot, Felix (2017). “Global production with export platforms”. In: *The Quarterly Journal of Economics* 132.1, pp. 157–209.
- World Bank (2008). *Global Purchasing Power Parities and Real Expenditures: 2005 International Comparison Program*. Washington, DC: World Bank. URL: <https://openknowledge.worldbank.org/handle/10986/21558>.
- (2024). *Viet Nam 2045: Trading Up in a Changing World — Pathways to a High-Income Future*. Tech. rep. P178784 / 10986/42505. © World Bank, published December 6, 2024. Handle: 10.1596/10986/42505. World Bank. URL: <https://documents1.worldbank.org/curated/en/099111424204523679/pdf/P1787841e077190d919b24181b4dcb14765.pdf>.

A Empirical Appendix

Figure B1: BTA Tariff Changes and Multinational Shares



Notes: This figure shows a binned scatterplot of the changes in tariffs due to the BTA against the foreign share. Industries are weighted by their share of Vietnam’s exports in 2000, and the solid line shows a weighted regression. The reported r is the weighted correlation computed on the underlying observations.

Table B1: Learning Spillovers from BTA and Multinationals on Foreign Firms

	# MNC Connected (1)	Tech Transfer (2)	R&D (3)	Modify Technology (4)	# Successes (5)
BTA	-11.31 (7.96)	-0.24*** (0.08)	-0.14*** (0.03)	0.10* (0.06)	0.49** (0.20)
Foreign Share	10.03*** (2.92)	-0.20*** (0.06)	0.07*** (0.02)	0.07** (0.03)	0.38*** (0.12)
Foreign Share \times BTA	-109.88** (49.19)	-0.14 (0.75)	1.20*** (0.17)	0.74* (0.40)	3.50** (1.49)
Controls	✓	✓	✓	✓	✓
Observations	303	303	377	377	377
Industries	79	79	79	79	79
R ²	0.39	0.19	0.25	0.13	0.14
Mean Dep. Var	8.57	0.27	0.09	0.07	0.15
KP Wald F	32.56	32.56	33.31	33.31	33.31

Notes: This table reports IV estimates for foreign-owned firms. Variable definitions, data sources, and estimation details are identical to Table 3.

B Theory Appendix

B.1 Hat Algebra

Theorem 1. *Assume we have data on a baseline static equilibrium, and consider a set of counterfactual changes to trade costs, multinational production costs, and technologies. Let a hat over a variable denote the proportional change in that variable between baseline and counterfactual equilibria. Then these proportional changes solve the following system of equations:*

$$\hat{x}_j^s = \hat{w}_j^{\gamma_j^s} \left(\Pi_r \left(\hat{P}_j^r \right)^{\beta_j^{sr}} \right)^{1-\gamma_j^s}, \quad (\text{H1})$$

$$\hat{\kappa}_{ijh}^s = \hat{x}_h^s \hat{\delta}_{jh}^s \hat{\tau}_{ij}^s, \quad (\text{H2})$$

$$\hat{K}_{ih}^s = \left(\sum_j \psi_{ijh}^s \left(\hat{\kappa}_{ijh}^s \right)^{\frac{-\theta}{1-\eta}} \right)^{\frac{1-\eta}{-\theta}}, \quad (\text{H3})$$

$$\hat{\psi}_{ijh}^s = \left(\frac{\hat{\kappa}_{ijh}^s}{\hat{K}_{ih}^s} \right)^{\frac{-\theta}{1-\eta}}, \quad (\text{H4})$$

$$\hat{\pi}_{ijh}^s = \hat{\psi}_{ijh}^s \hat{\lambda}_h^s \left(\frac{\hat{P}_i^s}{\hat{K}_{ih}^s} \right)^{\theta}, \quad (\text{H5})$$

$$\hat{P}_i^s = \left(\sum_h \pi_{ih} \hat{\lambda}_h^s \left(\hat{K}_{ih}^s \right)^{-\theta} \right)^{-\frac{1}{\theta}}, \quad (\text{H6})$$

where $\pi_{ih} = \sum_j \pi_{ijh}$ represents total sales from firms headquartered in h to buyers in i . Revenues, expenditures, and wages are given by:

$$\hat{R}_j^s R_j^s = \sum_h \sum_i \hat{\pi}_{jih}^s \hat{X}_i^s \pi_{jih}^s X_i^s, \quad (\text{H7})$$

$$\hat{X}_j^s X_j^s = \alpha_j^s \hat{w}_j w_j l_j + \sum_r (1 - \gamma_j^s) \beta_j^{rs} \hat{R}_j^r R_j^r, \quad (\text{H8})$$

$$\hat{w}_j w_j l_j = \sum_s \gamma_j^s \hat{R}_j^s R_j^s. \quad (\text{H9})$$

Additionally, suppose that the baseline static equilibrium is part of a balanced growth path (BGP), as defined in (...), in which technology in each country grows at rate λ_λ . Then, given a set of counterfactual changes to trade costs and multinational production costs, the law of motion for the proportional change in technology in h relative to the baseline at time zero is,

$$\hat{\lambda}_h^s(t) = \left(g_\lambda e^{(1-\rho)g_\lambda t} \right) \sum_i \sum_j a_{ihk}^s \hat{r}_{ihk}^s \left(\frac{\hat{\lambda}_k^s(t)}{\hat{\pi}_{ihk}^s(t) \hat{\psi}_{ihk}^s(t)^{-\eta}} \right)^\rho, \quad (\text{H10})$$

where

$$a_{ihk}^s = \frac{r_{ihk}^s \left(\frac{\lambda_k^s}{\pi_{ihk}^s \psi_{ihk}^{-\eta}} \right)^\rho}{\sum_{i'} \sum_{k'} r_{i'hk'}^s \left(\frac{\lambda_{k'}^s}{\pi_{i'hk'}^s \psi_{i'hk'}^{-\eta}} \right)^\rho}. \quad (\text{H11})$$

In the long run technologies will converge to their (counterfactual) BGP values, given by

$$\widehat{\lambda}_h^s(t) = \left(e^{(1-\rho)g\lambda t} \right) \sum_i \sum_j a_{ihk}^s \widehat{r}_{ihk}^s \left(\frac{\widehat{\lambda}_k^s(t)}{\widehat{\pi}_{ihk}^s \widehat{\psi}_{ihk}^{-\eta}} \right)^\rho. \quad (\text{H12})$$