The Impact of NAFTA on Prices and Competition: Evidence from Mexican Manufacturing Plants*

Ayumu Ken Kikkawa[a], Yuan Mei[b], Pablo Robles Santamarina[c]

[a]Sauder School of Business, University of British Columbia
[b]Singapore Management University
[c]The Brattle Group

This version: October 3, 2019

Abstract

This paper assesses the impact of the North American Free Trade Agreement on Mexican manufacturing plants' prices and markups. We distinguish between Mexican goods that are exported and those sold domestically, and decompose their prices separately into markups and marginal costs. We then analyze how these components were affected by reductions in Mexican output tariffs, intermediate input tariffs, and U.S. tariffs. We find that declines in these tariffs led to significant reductions in the marginal costs of Mexican products. However, prices of exported goods slightly increased as exporters increased their markups in response to declines in U.S. tariffs.

*First version: July 2018. Empirical analyses for this paper were conducted at the Microdata Laboratory Center of the Mexican Institute of Statistics and Geography (INEGI) in Mexico City. Conclusions and opinions expressed in this study are the sole responsibility of the authors and are not part of the official statistics of the National Statistical and Geographic Information System or INEGI. All results have been reviewed to ensure that no confidential information has been disclosed. We thank Rodrigo Adao, Jonathan Dingel, Akira Sasahara, Chad Syverson, and Felix Tintelnot for comments and discussions. Special thanks go to Liliana Martínez, Carmen Marquez, Andrea Ortiz, and Natalia Volkow from INEGI Mexico City, and Gabriel Romero from INEGI Aguascalientes, for their help with the data.

1 Introduction

The past few decades have seen drastic reductions in tariffs, a large share of which can be attributed to reciprocal multilateral or bilateral trade liberalizations. Reciprocal tariff reductions among participating countries may affect domestic firms' prices and the competition they face through multiple channels. As tariffs on foreign goods fall, domestic firms face more competition but at the same time can take advantage of cheaper imported inputs. In addition, as tariffs imposed on exported goods fall, exporters may enjoy greater access to foreign markets. It is therefore important to account for all of these channels together in order to better understand the economic effects of reciprocal trade liberalizations.

Despite the growing body of research on reciprocal trade liberalizations, their impact on prices and markups has yet to be empirically analyzed. Decomposing prices into costs and markups allows us to assess both the total gains from an episode of reciprocal trade liberalization and the distribution of gains between producers and consumers. However, it is empirically challenging to identify the impact of tariff reductions on markups and marginal costs because price data are not usually available at a disaggregated product-destination level for a broad set of industries. Moreover, since markups and marginal costs are not observable, a structural model must be used to estimate them.

In this paper, we empirically analyze the impact of reciprocal trade liberalizations by focusing on how Mexican firms responded to the North American Free Trade Agreement (NAFTA). To do so, we rely on a confidential dataset from Mexican manufacturing plants that includes disaggregated plant–product–level data for the period 1994–2008. The data record quantity and price (unit value) information for both domestic and exported products produced by plants that cover 85% of value added in manufacturing. A unique feature of this dataset is the distinction between domestic and foreign markets, which allows us to distinguish between the impact of reciprocal trade liberalization on exporters and the impact on domestic producers. Equipped with these data, we follow the empirical framework developed by de Loecker et al. (2016) and derive estimates of markups and marginal costs at the plant–product-destination level. This method estimates production functions to identify markups from the wedge between the output elasticity of a variable input and its expenditure share out of total revenue, which is now a standard approach in the industrial organization literature.[1] One advantage of this approach is that we do not need to make any assumptions about market structures or consumer preferences to recover markups.[2] In addition, we are able to estimate the product-level markups and marginal costs of multi-product firms across a broad set of manufacturing industries, which is a novel contribution to existing studies on NAFTA.[3]

Tariff reductions from NAFTA potentially affected markups and marginal costs of Mexican manufacturing plants via three different channels. The first is the *competition channel*, in which higher competitive pressure due to lower tariffs imposed by the Mexican authorities may have forced Mexican producers to lower their markups and compelled them to become more productive. The second channel is the *cost channel*: a decline in tariffs on intermediate inputs may have directly lowered the marginal costs for the Mexican plants importing those products. Lastly, since NAFTA was a reciprocal agreement, reductions in tariffs imposed by the United States may have enabled incumbent Mexican exporters to raise their factory-gate prices and markups without increasing the after-tariff price paid by U.S. consumers. We refer to this channel as the *market access channel*. We explore these three channels by examining the impact of Mexican output tariffs, tariffs on intermediate inputs, and U.S. tariffs on prices, markups, and marginal costs at the plant–product–destination level. For

[1]See Olley and Pakes (1996), Levinsohn and Petrin (2003), de Loecker (2011), de Loecker and Warzynski (2012), and Ackerberg et al. (2015) for production function estimation at the plant level. de Loecker et al. (2016) and Garcia-Marin and Voigtländer (2016) are examples of production function estimation at the product level.

[2]An alternative approach for markup estimation, exemplified by Berry et al. (1995), Goldberg (1995), and Goldberg and Hellerstein (2013), assumes specific preferences and market structure to derive estimates of markups. The detailed product-market-level data required as well as the particular assumption on market structure makes it impossible to use this approach for a broad set of industries as we do in this work.

[3]A large body of empirical studies has analyzed the impact of NAFTA on trade volume (Romalis, 2007), productivity (López-Córdova, 2003; Iacovone, 2012; de Hoyos and Iacovone, 2013), quality upgrading (Verhoogen, 2008), and income (Easterly et al., 2003; Esquivel and Rodríguez-López, 2003). We are not aware of any existing work that has studied NAFTA's impact on markups and marginal costs at the plant–product level.

each channel, we analyze both the average effects and the dynamic effects over time following Dix-Carneiro and Kovak (2017). In both cases, the impacts of the three channels are identified by exploiting the variation of these quantities within a plant–product–destination, controlling for changes in macroeconomic conditions at the sector–year level.

We first analyze the impact of NAFTA through the competition channel and the cost channel, both of which have been emphasized by de Loecker et al. (2016) and other existing works. For the competition channel, our results show that reductions in Mexican output tariffs had a significant effect on prices, markups, and marginal costs. For domestic products, we find that both marginal costs and prices decreased in response to reductions in output tariffs, but the changes in markups were insignificant. In addition, we also find evidence of pro-competitive effects of NAFTA: Once the impact on marginal costs is controlled for, output tariff reductions led to a decrease in markups. This effect is stronger for large firms and firms with higher initial markups. On the other hand, exporters, whose competition was not affected by the decline in Mexican output tariffs, responded to lower marginal costs by raising markups. The resulting effect on prices was insignificant. Over time, we find that marginal costs declined immediately after tariff cuts, but increases in markups were gradual.

For the cost channel, we find that exporters and non-exporters in Mexico responded differentially to the input tariff reductions. Input tariff reductions had no significant effect on the markups and marginal costs of domestic products on average. Over time, input tariff reductions on domestic products led to a gradual increase in marginal costs after an initial drop. To rationalize this dynamic pattern, we construct proxies of quality following the intuition introduced in Khandelwal (2010) and find supporting evidence of quality upgrading over time. On the other hand, declines in input tariffs significantly reduced the marginal costs of exported products. Exporters, who were more likely to import inputs from abroad, increased markups in response to the reduction in marginal costs, resulting in a zero effect on the prices.

In addition to the competition channel and the cost channel, we also find significant evidence of the impact of NAFTA through the market access channel. This channel has not been studied extensively because plant–product-level data that distinguish between exporters and domestic producers are rare. We find that declines in U.S. tariffs on exported Mexican products led to a substantial reduction in the marginal costs as Mexican exporters increased their productivity. This decrease in the marginal costs is consistent with the finding in Bustos (2011) that Argentinian exporters increased their spending on technology in response to tariff reductions by Brazil. We also find that markups on exported products increased by more than the decline in marginal costs, leading to a slight increase in the prices on exported products. This increase in markups is observed even after controlling for marginal costs, suggesting that exporters responded to declines in U.S. tariffs on Mexican products by raising their markups. Moreover, this effect is stronger for smaller exporters. We do not find any existing empirical study documenting this kind of anti-competitive effect during an episode of reciprocal trade liberalization.

Overall, our estimates imply that the observed tariff declines during the 1994–2008 period led to an average reduction in prices and marginal costs of Mexican domestic products by 3.4% and 8.2%, respectively. At the same time, markups increased by 4.8% on average. For products exported to the United States, the tariff reductions led to a small price increase of 1.9%, a reduction in marginal costs of 27.9%, and an increase in markups of 29.9%. These results suggest that Mexican consumers benefited from NAFTA through lower prices. On the other hand, Mexican producers enjoyed more profits through lower input prices and higher markups.

Our work complements existing research on estimating the competitive effects of a trade liberalization by providing the first assessment of the impact of a free trade agreement on both prices and markups.[4] In this strand of literature, our work is most closely related to de Loecker et al. (2016) who estimate product-level markups and analyze trade liberalization

[4] Viewing free trade agreements (FTAs) from the demand side, Faber (2014) uses microdata underlying Mexican consumer price index and analyzes how import tariff cuts through NAFTA affected product quality and income inequality.

in India in the 1990s.[5] Relying on the same estimation framework, both our paper and de Loecker et al. (2016) find pro-competitive effects from output tariff declines and the incomplete cost pass-through to prices. In addition to the competition channel and the cost channel studied in de Loecker et al. (2016), we also find evidence of anti-competitive effects from improved foreign market access for Mexican exporters. This market access channel distinguishes between reciprocal and unilateral trade liberalizations, but existing studies such as Lileeva and Trefler (2010) and Bustos (2011) mostly focus on the impact of improved market access on productivity instead of markups.

Our second contribution is to separately identify the effect of tariff reductions on goods that are exported and sold domestically, compared to existing studies that do not distinguish between foreign and domestic markets.[6] For example, both de Loecker et al. (2016) and Brandt et al. (2017) find a strong negative relationship between input tariffs and markups. Our results for Mexico confirm this negative relationship among the exported goods but not among those serving the domestic market. In other words, the benefit of input cost reduction during episodes of trade liberalization emphasized by de Loecker et al. (2016) and previously Amiti and Konings (2007) only applies to Mexican exporters. We find this pattern to be intuitive because only a small share of non-exporters in Mexico use imported intermediates.

The empirical findings presented in this paper are in line with the theoretical literature that analyzes the competitive effects of trade reforms. Within this literature, our work is closely related to the trade model analyzed by de Blas and Russ (2015) in which the endogenous distribution of markups responds to changes in trade costs.[7] By lowering trade costs,

[5]Other contributions to this literature are the studies by Levinsohn (1993) on Turkey, Harrison (1994) on the Ivory Coast, Krishna and Mitra (1998) on India, and Brandt et al. (2017) on China. Caselli et al. (2017) also follow the same method as de Loecker et al. (2016) to estimate markups of Mexican plants at product level, but the focus of that paper is on exchange rate pass-through.

[6]One exception is the study by Blum et al. (2018), which estimates production technologies and markups of Chilean manufacturing firms for each destination market in order to disentangle demand and cost drivers of firm-heterogeneity.

[7]Quantitative trade models, like that from Eaton and Kortum (2002) with perfect competition, and the monopolistic competition model of Melitz (2003), are unable to capture the competitive effects of trade liberalization since they assume constant markups. Even in models with variable markups, few predict changes to competition from liberalization. Bernard et al. (2003) and Arkolakis et al. (2017), for example, allow for variable markups and find that the distribution of markups is invariant to changes in trade costs.

trade liberalization indirectly reduces the residual demand for domestic goods, leading to a decline in domestic markups (pro-competitive effects) and an increase in welfare. However, in a regional free trade agreement, the increase in welfare is offset by a rise in foreign markups (the anti-competitive effects in de Blas and Russ, 2015), which lowers the overall gains from trade. Our paper contributes to this literature by providing the first empirical evidence for both pro-competitive and anti-competitive effects of reciprocal trade liberalizations.

Finally, our paper is related to the literature examining the effects of trade liberalization on productivity, which mostly emphasizes two main channels. The first channel is through foreign competition, as increased competitive pressure compels producers to be more efficient.[8] The second channel is through cheaper imported intermediate inputs, which not only lower the cost of production, but also allow for substantial technological improvements by expanding the set of available inputs.[9] Within this literature, our work is closest to the papers of López-Córdova (2003) and Iacovone (2012), which analyze the impact of tariff declines during NAFTA on plant-level measures of revenue productivity for Mexican manufacturing plants. One potential issue of using revenue productivity is biased estimates: If more efficient producers charge lower prices, this will be reflected as lower revenue productivity.[10] Our paper, by contrast, analyzes the impact of tariff reductions on marginal costs and hence is immune to this problem.

The rest of the paper is structured as follows. Section 2 briefly discusses the background of NAFTA and the data used in the estimation, before performing a preliminary analysis on the impact of tariff reductions from NAFTA on prices. Section 3 formally presents the empirical framework used in the estimation of markups and marginal costs. Section 4 shows the estimates of markups and marginal costs and provides evidence on the validity of the

[8] Examples of papers that emphasize this channel include Trefler (2004) and Lileeva and Trefler (2010) for the Canada-U.S. Free Trade Agreement, and Pavcnik (2002), who studies trade liberalization in Chile.

[9] The analyses of Amiti and Konings (2007) for Indonesia and Khandelwal and Topalova (2011) for India emphasize this second channel.

[10] For example, Pierce (2011) documents different effects of antidumping duties on the physical and revenue productivity measures of U.S. manufacturing plants. See Garcia-Marin and Voigtländer (2016) for a detailed discussion on this point.

estimation results. Section 5 analyzes the impact of tariff reductions on prices, markups, and marginal costs, and the last section concludes.

2 Background, Data, and Impact of NAFTA on Prices

In this section, we first briefly discuss NAFTA and the role of Mexico when this agreement was negotiated. We also describe the data used in this paper and provide some preliminary empirical facts on the effect of tariff reductions on prices, before formally introducing our estimation framework in Section 3.

2.1 Mexico and NAFTA

The Mexican experiment with trade liberalization started in the late 1980s. Mexico joined the General Agreement on Tariffs and Trade (GATT) in 1986, agreeing to a 50% ceiling on its tariffs. A subsequent liberalization in 1987 reduced and homogenized tariffs to 0%, 5%, 10%, or 20%, eliminated quotas, and reduced the number of products covered by official import prices. Although these advances represented a significant break from Mexico's previous protectionist policies, significant barriers to trade remained. In 1993, for example, only 10% of manufacturing products had tariffs at or below 15%.[11]

However, when Carlos Salinas began his presidency in 1988, further liberalization was not his priority. After nearly a decade of recession, high inflation, and a significant debt burden, Salinas and his team of technocrats wanted to implement structural reforms to modernize the Mexican economy. A key element of these structural reforms was an increase in investment in multiple sectors. Unfortunately, the previous decade had left the Federal Government without the funds needed to undertake the required investments. Also, prior debt binges and forced restructuring had locked Mexico out of international debt markets. The solution, the government decided, was to attract foreign direct investment. In January 1990, the need

[11] See Ten-Kate (1992) and López-Córdova (2003).

for foreign direct investment led the president and members of his cabinet to an unsuccessful roadshow across Europe. The fear of arbitrary expropriation by the Mexican authorities, as well as the integration of former communist countries into the European market, reduced the relative attractiveness of Mexico as an investment destination. Running out of options, the Mexican team tried one last gambit. During the final leg of the trip at the World Economic Forum in Davos, Switzerland, President Salinas approached the U.S. Trade Representative, Carla Hills, with the idea of creating a single market between Mexico and the United States. The hope was that a single market would attract the foreign direct investment that further unilateral liberalization could not. Three years and eleven months later, after a period of intense negotiations, the North American Free Trade Agreement between Canada, Mexico, and the United States went into effect.

The urgency to reach an agreement by a Mexican government starved for funds led Mexican negotiators to concede to nearly every pressure from the dominant U.S. and, to a lesser extent, Canadian negotiators. Cameron and Tomlin (2002)'s account of negotiations in Dallas in 1992 illustrates this point. They document that during those talks, Chief Mexican Negotiator Herminio Blanco wrote a memorandum to the Mexican negotiation team with a message along the lines of "show your cards, get to the bottom, there is no tomorrow." The ensuing enthusiasm with which Mexican negotiators made concessions surprised the Canadian and U.S. negotiators. This evidence, coupled with the fact that many provisions included in NAFTA came from the previous Canada-U.S. Free Trade Agreement (CUSFTA), suggests that NAFTA provisions can plausibly be considered exogenous to the interests of Mexican industries, and more importantly for our purposes, to the interests of individual Mexican plants and product lines.

NAFTA was a comprehensive agreement that covered a range of issues, including market access, intellectual property, and investment protection provisions. We focus on the impact of one aspect of the agreement: the reductions in tariffs. Under NAFTA, the Governments of Canada, Mexico and the United States agreed to eliminate tariffs on most of their products by the year 2008. Tariffs were scheduled to decline annually at a constant rate until they were completely phased out. However, there was substantial heterogeneity in the actual phase-out dates. For some products, the phase-out started in 1994, the year that NAFTA came into effect. The phase-out years for other products were 1998, 2003, and 2008, respectively. Figure 1 shows the Mexican tariff schedule for a selected sample of products. We can see that the heterogeneity in phase-out dates and in the initial level of tariffs before NAFTA led to substantial variation in yearly tariff cuts across products.[12]

2.2 Data

2.2.1 Manufacturing Survey Data

We use manufacturing plant and product data from two surveys conducted and maintained by the Mexican Institute of Statistics and Geography (INEGI): the Monthly Industrial Survey (EIM) and the Annual Industrial Survey (EIA), available for the 1994–2008 period. Both surveys initially cover the same set of plants based on the 1993 Economic Census. The EIA, however, was updated in 2003 to include new plants that were either opened after 1994 or identified by INEGI after that year. Since our analysis requires information on variables from both the EIA and EIM, we include only plants present in both the EIM and EIA in the estimation. These surveys classify plants according to a unique 6-digit class (similar to the 6-digit North American Industrial Classification System (NAICS) industry code) based on the *1994 Mexican Classification of Activities and Products (CMAP94)*, a precursor of NAICS. The surveys cover 206 6-digit classes in the manufacturing sector. The number of plants included was chosen to ensure that the sample covered at least 85% of the value added in each class and contained all plants with more than 100 employees. The final sample of plants represents around 85% of the value added in manufacturing in Mexico. For a more detailed description of the sampling methodology and data, refer to Online Appendix Section 1.

[12] Besedes et al. (2018) find no evidence that this heterogeneity in phase-out schedules impacted U.S. import growths from Mexico across products.

These surveys, however, have some limitations. Although coverage concerning value added is excellent, the sample is skewed toward larger plants. Furthermore, while the surveys track plant exit systematically, they do not record plant entry.[13]

The EIM survey reports monthly data on employment and the wage bill at the plant level, and quantities and sales value of disaggregated products separated by products destined for the domestic and those earmarked for the export market. The distinction of products between domestic and export markets is a unique feature of the EIM. The data do not include the destinations of exports. However, Mexico's exports are highly concentrated, with more than 85% of exports going to the United States during the period examined. Since the majority of these products were going to the United States, we assume that all exported products were destined for that country. Products are disaggregated at the 8-digit level and can be viewed as individual product lines. See Online Appendix Section 1 for examples of these product lines. These products are very disaggregated relative to other manufacturing surveys that include product or product line information. The availability of sales and quantity data by product enables construction of unit values as a measure of prices. We aggregate monthly values into annual data to match them with information from the EIA.

The EIA survey records plant-level yearly information on plant inputs, total production, and other detailed data on plant operations. Most of the manufacturing plant data used in our estimation, except for quantities and sales at the product level, come from this survey. We use the survey data on material expenditures, total employment, capital, import and export status, research and development expenditures, and plant's state.[14] We deflate all monetary variables by their appropriate price deflator. For more information on the construction of the variables and other sources of external data used in the estimation, see Online Appendix Section 1.

[13]The survey also lacks information on certain dimensions that may be important in Mexican exports. First, it does not cover the so-called Maquiladora plants that mostly engage in processing exports specialized in labor-intensive products (Utar and Ruiz, 2013). Second, it does not record whether the plants are part of multinational enterprises.

[14]We construct the capital series with the perpetual inventory method using investment by type and the initial book value of a capital stock. See Online Appendix Section 1 for more details.

We match the EIM and EIA surveys using a unique plant identifier provided by INEGI and construct a panel of approximately 180,000 product–plant–year observations from 1994 to 2008. Table 1 shows the average number of plant–product–destination observations by sector, as well as the average number of products by a plant in the sample. Table 2 presents the number of plants in the sample, as well as summary statistics of the main variables from the EIA that we use in the estimation by sector. We can see that the majority of plants in the sample are multi-product and non-exporter plants. However, the number of single-product plants per sector in the sample is relatively large. This feature contributes to the empirical strategy discussed in later sections.

2.2.2 Tariff Data

Tariff data for Mexico and the United States come from the World Bank's *World Integrated Trade Solution*, available at the HS 6-digit level. Since there is no concordance available between HS-6 classification and the Mexican *CMAP94* classification, we manually constructed the concordance. The process involved matching approximately 5,000 products in the *CMAP94* classification to one or sometimes multiple HS codes using the *CMAP94* product description provided by INEGI. When multiple HS-6 codes corresponded to a single *CMAP94* product, we used the average tariff across corresponding HS-6 codes. We explain the details in Online Appendix Section 6. Utilizing this concordance, we constructed a measure of output tariffs, i.e., the tariffs applied by the Mexican authorities to goods coming from the United States. We also constructed a measure of U.S. tariffs applied to goods coming from Mexico using the same concordance.

To capture the impact of cheaper access to imported intermediates that resulted from tariff declines under NAFTA, we constructed class-level intermediate input tariffs using the Mexican input–output tables available from INEGI for 2003 at the 4-digit NAICS classification and matched them to a *CMAP94* classification using the concordances provided by

INEGI.[15] To do this, we calculated the average output tariff within a class and then used direct and indirect requirement coefficients from the I-O table to compute a weighted-average tariff at the class level. Formally, the intermediate input tariffs of plants from class j at time t is given by:

$$\tau_{jt}^{input} = \sum_{k} \Phi_{kj} \cdot \tau_{kt}^{output},$$

where Φ_{kj} is class j's share of intermediate inputs coming from class k, and τ_{kt}^{output} is the average output tariff in class k. See Online Appendix Section 1.2 for more details on the construction of the tariff measures.

To summarize, the output and U.S. tariffs used in our analysis vary at the product level, whereas intermediate input tariffs vary at the class level. Table 3 gives summary statistics on the tariff measures for the year 1993 (the year before NAFTA started) and 2008 (the last year in the sample). As we can see, output tariffs declined by 14.6 percentage points on average during this period. The measure of intermediate input tariffs decreased by 9.3 percentage points. Finally, U.S. tariffs on Mexican products fell by only 5.1 percentage points, as their level before NAFTA was lower than the initial level of Mexican tariffs. Figure 2 shows the average of our three measures of tariffs for the period 1993–2008.

2.3 The Impact of Tariff Declines on Prices

Before decomposing prices into markups and marginal costs, we examine the impact that declines in our tariff measures had on prices. We estimate the following specification separately for domestic and exported products:

$$\log P_{ijt} = \alpha + \beta_1 \tau_{it}^{output} + \beta_2 \tau_{jt}^{input} + \beta_3 \tau_{it}^{US} + \xi_{ij} + \psi_{st} + \varepsilon_{ijt}, \tag{1}$$

[15] This is the only I-O table available in Mexico with a high degree of disaggregation, since I-O tables before 2003 are available only at a highly aggregated level.

where P_{ijt} is the price of product i from plant j at time t, τ_{it}^{output} is the Mexican output tariff applied to product i, τ_{jt}^{input} is the intermediate input tariff of plant j, τ_{it}^{US} is the tariff applied by the United States on good i from Mexico, and ξ_{ij} and ψ_{st} are plant–product and sector–year fixed effects, respectively. Since the primary policy variable, output tariffs, varies at the product level, we cluster standard errors at the product level. As tariffs enter in levels in equation (1) and prices in logarithms, the coefficients βs are semi-elasticities. That is, they measure the percent change in prices when tariffs increase by one percentage point. We include U.S. tariffs on Mexican products in the specification for two reasons. First, since most exported products went to the United States, changes in U.S. tariffs should have affected their prices. Second, declines in U.S. tariffs might have influenced the behavior of Mexican plants, affecting both exported and domestic products. Iacovone and Javorcik (2012), for example, find that an increase in market access driven by a decline in U.S. tariffs has stimulated investment by Mexican manufacturing plants as they prepare to introduce new products into the export market.[16]

The coefficients for the tariffs are identified by exploiting variation in the dependent variable and tariffs within a plant–product–destination over time, controlling for changes in macroeconomic conditions at the sectoral level.

2.3.1 More on Identification: Exogeneity of Tariff Cuts

The main concern with this specification is that tariff changes might be correlated with omitted factors that also affect Mexican sectoral outcomes. However, by focusing on the disaggregated product level, identification requires the more plausible assumption that tariff cuts are exogenous at the level of individual plant–product pairs. In this particular exercise, the main concern for identification would be if tariff schedules under NAFTA were set to protect specific products or industries. If this protectionism was a factor in determining tariff schedules, the previous specification would suffer from endogeneity problems, as tariffs

[16]See also Head and Ries (1999), in which they explore the effect of CUSFTA on Canadian firms' productivity through increases in scale.

would be correlated with an omitted factor that would also determine prices: protectionism. If this were the case, we would be overestimating, for example, the impact of output tariffs on prices.[17]

There is substantial evidence, both empirical and anecdotal, that endogeneity arising from protectionism is not likely to be the case. If tariffs were set with protectionism in mind, we would expect that products with high initial tariffs would face higher tariffs under NAFTA than the average product, and be subject to a slower tariff decline schedule. The data show, however, that products with high initial tariffs faced the largest tariff declines from NAFTA (See Online Appendix Section 1.2). Moreover, Kowalczyk and Davis (1998) present empirical evidence that indicates that phase-out periods for Mexican tariffs appear to be uncorrelated with their levels before NAFTA. From the anecdotal side, as discussed in Section 2.1, the circumstances surrounding NAFTA negotiations suggest that Mexican negotiators had very little bargaining power in setting tariffs. Both sets of evidence suggest that, during the analysis, we can plausibly consider tariff reductions under NAFTA to be exogenous from the viewpoint of individual Mexican plant–product pairs.[18] We, therefore, maintain the assumption that tariff cuts from NAFTA are exogenous for the rest of the paper.

Another potential identification problem is that the three measures of tariffs are likely to be correlated, so there might not be enough variation in the data to identify the coefficients in equation (1) separately. Table 7 in the Online Appendix shows that the three tariff measures are indeed correlated. However, the correlation is far from perfect, suggesting that there is enough variation left over to identify the coefficients separately.

[17] See Online Appendix Section 2 for a formal argument.

[18] This assumption also implies that the tariff reductions were orthogonal to the extent that plants faced different treatments indirectly through their domestic suppliers and buyers. Moreover, Conconi et al. (2018) find that NAFTA Rules of Origin reduced the growth of Mexican imports of intermediates from third countries relative to the United States and Canada. Our analysis of Mexico is to a large extent immune to the potential endogeneity of NAFTA Rules of Origin, as around 90% of the NAFTA rules were already present in CUSFTA.

2.3.2 Results

Table 4 presents the results of the estimation of equation (1). We can see that the coefficient for output tariffs of domestic products is positive and statistically significant. This result is in line with the argument that increased foreign competition places competitive pressures on domestic producers, pushing them to reduce prices. The magnitude of the coefficient means that a one percentage point decline in tariffs led to a 0.13% decline in prices. Coefficients for input and U.S. tariffs on domestic products are not statistically significant. Using the estimated coefficients and the observed average tariff declines from 1993, before the beginning of NAFTA, to 2008, we find that tariff declines led to a statistically significant 3.43% reduction in the prices of domestic products.

For exported products, the effect of declines in output tariffs or intermediate input tariffs on prices is not statistically significant. We find that U.S. tariffs had a statistically and economically significant impact on the prices of exported products, with estimates suggesting that a one percentage point decline in U.S. tariffs led to a 0.50% increase in prices. Therefore, it appears that exporters responded to the reductions in U.S. tariffs by raising prices, partially offsetting the decrease in tariffs. Prices, however, increased by less than the decline in U.S. tariffs. In other words, the after-tariff price paid by U.S. consumers continued to decline. Finally, the observed decrease in tariffs seems to have raised the price of exported products slightly, but this increase is not statistically significant.

To better understand the forces driving the results in Table 4, we analyze the impact of tariffs on prices through their effect on marginal costs and markups. In the following section, we present the empirical framework used to decompose observed product-level prices into their unobserved markup and marginal cost components.

3 Empirical Framework

In this section, we set up and discuss the empirical framework used to recover markups and marginal costs at the product level, separately for domestic and exported products. We use the framework developed by de Loecker et al. (2016), which relies on the estimation of quantity production functions and exploits plants' first-order conditions to arrive at estimates of markups and marginal costs at the plant–product level.[19]

Estimating production functions at the product level offers many advantages over standard plant-level estimation, but it also poses some important empirical challenges since the multi-product feature of plants must be taken into consideration. In particular, it forces us to deal with the fact that we do not observe input allocations at the product level in the data.[20] We can surmount this challenge through some assumptions, as demonstrated below. Essential to resolving this issue is the suggestion from de Loecker et al. (2016) that a researcher can estimate production functions using a sample of single-product plants, and then recover the unobserved input allocations for multi-product plants from restrictions imposed by the structural model. We follow methods from de Loecker et al. (2016) closely, with slight modification to the product definition in order to account for the unique features of the data.

As mentioned in Section 2.2.1, the data include quantities and sales revenues separately for domestic and exported products. One issue is how to treat these products during estimation. We assume that domestic and exported varieties of the same disaggregated product category are distinct products, even if the same plant produced them. We treat, for example, bicycles destined for the domestic market as a different product from bicycles intended for the export market. We treat domestic and exported varieties as different products for two reasons. First, products destined for the domestic market may differ in quality from exported products from the same plant. If quality is positively correlated with consumer income, for

[19]For a more general exposition of the framework, see Online Appendix Section 2 and de Loecker et al. (2016).

[20]To our knowledge, Cajal-Grossi et al. (2019) is the only paper that directly observes input allocations at this level of disaggregation. However, they focus on the garment exporters in Bangladesh, instead of all manufacturing sectors.

example, we would expect the products exported to the United States to be of higher quality than those sold to the domestic market.[21] Second, since domestic and exported products are shipped to two different markets, they might respond very differently to changes in tariffs, even if they are exactly the same physical product. This treatment is consistent with the model and evidence presented by Verhoogen (2008), in which a single plant manufactures products of varying qualities for different markets. Therefore, in what follows, we treat each product–market pair as a distinct product. We describe the main elements of the framework below.

3.1 Recovering Markups and Marginal Costs

We use a structural model of production to obtain estimates of markups and marginal costs at the product level. Consider the following production function of product–market i from plant j in sector s at time t:

$$Q_{ijt} = F_i(M_{ijt}, L_{ijt}, K_{ijt}; \beta_s)\Omega_{jt}, \tag{2}$$

where Q_{ijt} is the physical output, M_{ijt} is the material input, L_{ijt} denotes the labor input, K_{ijt} is the capital input, β_s is the parameter vector of the production function that we assume is sector-specific, and Ω_{jt} is the Hicks-neutral productivity at the plant level. We restrict the parameters of the production functions to be the same for all products within a sector.

To estimate equation (2), we must make some additional assumptions. We assume that production function $F_i(\cdot; \beta_s)$ is product–market specific and twice differentiable with respect to materials, M_{ijt}. This assumption implies that a multi-product plant that produces SUVs and sedans, for example, uses the same production technology in the production of SUVs as a single-product plant that manufactures SUVs only. Their productivities are allowed to differ, however. We also assume that all of the inputs used by a plant are allocated to the

[21] See, for example, Linder (1961), Hallak (2006), and Hallak and Sivadasan (2013).

production process and that the share of input expenditures by product is the same across all inputs.[22] Finally, we assume that plants minimize short-run costs, taking quantities and input prices of materials, labor, and capital as givens.

We follow the approach of Hall (1986), subsequently refined by de Loecker and Warzynski (2012) and de Loecker et al. (2016), using a plant's first-order conditions from its cost minimization problem to recover the model's implied markups. Consider the Lagrangian of the cost minimization problem of product–market i from plant j at time t:

$$\begin{aligned}\mathcal{L}(M_{ijt}, L_{ijt}, K_{ijt}, \lambda_{ijt}) &= W^M_{ijt}M_{ijt} + W^L_{ijt}L_{ijt} + W^K_{ijt}K_{ijt} \\ &+ \lambda_{ijt}[\bar{Q}_{ijt} - Q_{ijt}(M_{ijt}, L_{ijt}, K_{ijt}, \Omega_{jt})],\end{aligned}$$

where $Q_{ijt}(M_{ijt}, L_{ijt}, K_{ijt}, \Omega_{jt}) = F_i(M_{ijt}, L_{ijt}, K_{ijt}; \beta_s)\Omega_{jt}$, λ_{ijt} is the Lagrange multiplier from the cost minimization problem, $\bar{Q}_{ijt}$ is a given level of output, and W^M_{ijt}, W^L_{ijt}, and W^K_{ijt} are the input prices of materials, labor, and capital, respectively. Taking the first-order condition of the Lagrangian with respect to materials, we have:

$$\frac{\partial \mathcal{L}_{ijt}}{\partial M_{ijt}}: \quad W^M_{ijt} - \lambda_{ijt}\frac{\partial Q_{ijt}}{\partial M_{ijt}} = 0, \tag{3}$$

where λ_{ijt} in equation (3) is the marginal cost of production.[23] Using this fact, we can rewrite the first-order condition to obtain an expression for markups. In particular, let μ_{ijt} denote the markup of product–market i from plant j. We have:

$$\mu_{ijt} = \theta^M_{ijt} \times (\Psi^M_{ijt})^{-1}, \tag{4}$$

where $\theta^M_{ijt} = \frac{\partial Q_{ijt}}{\partial M_{ijt}}\frac{M_{ijt}}{Q_{ijt}}$ is the output elasticity of material inputs, $\Psi^M_{ijt} = \frac{W^M_{ijt}M_{ijt}}{P_{ijt}Q_{ijt}}$ is the

[22] This means, for example, that if a product uses 20% of plant-level material expenditures, it must also use 20% of labor and 20% of capital. While restrictive, it allows us to identify input allocations and productivities for multi-product plants.

[23] The Lagrange multiplier tells us how much the objective function changes if we relax the constraint by one unit. In the cost minimization case, it tells us the extent to which costs increase if we increase production by one unit, which is precisely the definition of marginal costs.

expenditure share of materials in product i's revenues, and P_{ijt} is its sales price.[24]

The derivation of equation (4) highlights the need to have at least one input that is freely adjustable, since equation (4) comes from the static, first-order condition of a plant's cost minimization problem. We use materials as the flexible input. Although, in principle, we could use labor as the flexible input, as de Loecker and Warzynski (2012), using labor in a Mexican context may be problematic. Juarez and de la Cabada (2016) find significant labor market rigidities in Mexico from 1996 to 2011 that indicate adjustment costs and make labor a dynamic input, thus violating the validity of the first-order condition.

Equation (4) reveals that we need estimates of output elasticity, product revenue, and input expenditures per product in order to construct markups at the product level. Unfortunately, the Mexican data do not contain information on input expenditures by product in a plant. For example, the data include total expenditures on labor, capital, and materials for a producer of bicycles and tricycles, but not what portions were allocated to the production of bicycles and tricycles separately. Moreover, under general production functions such as the translog production function we use in the estimation, output elasticities will depend on these unobserved input expenditures. One exception is the Cobb-Douglas production function. Under the assumptions described above, output elasticities of Cobb-Douglas production functions are constant across products in a sector. The downside is that all variation in product markups comes from differences in input expenditure shares. Because of this, we have chosen the more flexible translog specification.

For these reasons, we take additional steps to recover input expenditures by product for multi-product plants. We follow the procedure developed by de Loecker et al. (2016). This procedure involves estimating the parameters of $F_i(\cdot; \beta_s)$ using single-product plants and employing these estimates and the equations implied by the structural model to recover these unobserved input expenditures. We develop this procedure in more detail below. Once

[24] Under perfect competition, output elasticity equals the expenditure share, and the markup is therefore equal to one.

we estimate the markup, we can obtain estimates of a product's marginal cost, MC_{ijt}, from:

$$MC_{ijt} = \frac{P_{ijt}}{\mu_{ijt}}.$$

Thus, once we estimate the markups, we have all of the elements needed to decompose the observed prices into unobserved markups and marginal costs.

3.2 Estimation of the Production Function

Allowing for log-additive measurement errors in output, we have the following equation by taking the logs of equation (2):

$$q_{ijt} = f_i(\mathbf{x_{ijt}}; \beta_\mathbf{s}) + \omega_{jt} + \epsilon_{ijt}. \tag{5}$$

$\mathbf{x_{ijt}} = (m_{ijt},\ l_{ijt},\ k_{ijt})$ is the log of inputs, q_{ijt} is the log of output, ω_{jt} is the log of productivity, and ϵ_{ijt} is the measurement error. Following the literature, we estimate a translog production function on materials, labor, and capital. The translog production function of product–market i of plant j in sector s is:

$$\begin{aligned} f_i(m_{ijt}, l_{ijt}, k_{ijt}; \beta_s) =& \beta_{sm} m_{ijt} + \beta_{smm} m_{ijt}^2 + \beta_{sl} l_{ijt} + \beta_{sll} l_{ijt}^2 + \beta_{sk} k_{ijt} + \beta_{skk} k_{ijt}^2 \\ & + \beta_{sml} m_{ijt} l_{ijt} + \beta_{smk} m_{ijt} k_{ijt} + \beta_{slk} l_{ijt} k_{ijt} + \beta_{smlk} m_{ijt} l_{ijt} k_{ijt}, \end{aligned} \tag{6}$$

where m_{ijt}, l_{ijt}, and k_{ijt} represent the logarithm of materials, labor, and capital, respectively, for plant j used in the production of product–market i at time t. One feature of translog production functions is that elasticities vary not only at the sector level, but also at the product–market level within a plant. Therefore, a plant that manufactures products in the same category for export and domestic markets has different output elasticities for each market. By not restricting elasticities to be the same across plants and products within a sector, as is the case with a Cobb-Douglas production function, the translog specification does not limit variation in markups to come exclusively from heterogeneity in input expenditure shares from sales revenue. To estimate equation (5), however, we must resolve some issues.

The presence of unobserved productivity ω_{jt} in equation (5) leads to simultaneity bias. This bias arises because plants observe their productivity draws before making their choice of inputs. Not taking this bias into account during the estimation yields inconsistent estimates of the parameters. If we had data on physical inputs at the product level, we could solve these problems following the proxy methods used by Olley and Pakes (1996) and Levinsohn and Petrin (2003). Unfortunately, we do not have input allocations across products in a plant.

An additional obstacle is the lack of data on physical units of inputs. Instead, we have data on input expenditures deflated by industry-level input price indices. Failure to observe input prices at the plant or product level when estimating quantity production functions might lead to significant biases in the estimation.[25] To understand how the results can be biased if input expenditures are used directly in the estimation, consider the following example. Suppose two Mexican producers of bicycles manufacture their products using materials as their only input. Assume that the first producer makes its bicycles using cheap domestic materials, and the second uses more expensive imported materials. Given their cost structures, suppose the first produces twice as many bicycles, spending half of what the second spends, but the second producer has larger revenues due to its higher quality product. Naive application of OLS for the number of bicycles sold on material expenditures deflated by a common industry-specific price index would yield a negative output elasticity of materials. Failure to correct for input price differences, therefore, could also lead to biased estimates.

We introduce some additional notation in order to further elucidate the econometric problems. Let ρ_{ijt} denote the logarithm of the share of input expenditures for product–market i of plant j assumed to be the same across inputs. Since all of the inputs are

[25] See the argument by de Loecker and Goldberg (2014).

assumed to be allocated to the production of products, the ρ_{ijt}s of plant j have to satisfy:

$$\sum_{i \in J_j} \exp(\rho_{ijt}) = 1,$$

where J_j is the set of products produced by plant j. Although in principle we would like to have real input allocations at the product level, $\mathbf{x_{ijt}} = (m_{ijt}, l_{ijt}, k_{ijt})$, in practice we have plant-level input expenditures that are deflated by industry-level input price indices, $\hat{\mathbf{x}}_{\mathbf{jt}} = (\hat{m}_{jt}, \hat{l}_{jt}, \hat{k}_{jt})$. Let $\mathbf{w_{ijt}}$ denote the vector of unobserved deviations of the logarithm of product i's input prices from the log of the industry-specific input price indices. We can write the vector of inputs $\mathbf{x_{ijt}}$ of product i from plant j as:

$$\mathbf{x_{ijt}} = \rho_{ijt} + \hat{\mathbf{x}}_{\mathbf{jt}} - \mathbf{w_{ijt}}.$$

Substituting this expression into equation (5), we have:

$$q_{ijt} = f_i(\hat{\mathbf{x}}_{\mathbf{jt}}; \beta_{\mathbf{s}}) + \Gamma(\rho_{ijt}, \hat{\mathbf{x}}_{\mathbf{jt}}; \beta_{\mathbf{s}}) + \Lambda(\mathbf{w_{ijt}}, \rho_{ijt}, \hat{\mathbf{x}}_{\mathbf{jt}}; \beta_{\mathbf{s}}) + \omega_{jt} + \epsilon_{ijt}, \tag{7}$$

where $\Gamma(\cdot)$ comes from the fact that we do not have product-level input allocations, and $\Lambda(\cdot)$ from the fact that we do not have plant–product-level input prices. From equation (7), it is clear that even if we control for unobserved productivity, ω_{jt}, estimates of β_s would still be biased, since both $\Gamma(\cdot)$ and $\Lambda(\cdot)$ depend on the vector of input expenditures, $\hat{\mathbf{x}}_{\mathbf{jt}}$. Estimation of the production function, therefore, requires us to deal with several issues: (1) the unobserved input expenditures by product, (2) the unobserved product-level input prices, and (3) the unobserved plant-level productivity. We address each of these below.

Unobserved Input Expenditures by Product

The assumptions above imply that multi-product plants and single-product plants that manufacture the same product use the same technology. For example, a plant that produces bicycles and tricycles uses the same technology to produce bicycles as a plant that produces bicycles only. This observation suggests that we can use single-product plants to estimate the parameters of the production functions. For single-product plants, the input expenditure share is one; therefore $\rho_{ijt} = 0$, and term $\Gamma(\cdot)$ in equation (7) is zero. This strategy, however, leads to a new type of sample selection bias, since we would be conditioning on plants that produce only a single product. Moreover, there would be another layer of selection bias, since most of the single-product plants only serve the domestic market. To correct for these biases, we use an unbalanced panel of plants that produce a single product serving either the domestic or foreign markets at any point in time. Thus, our estimating sample includes plants that produce a single product for a single market in a given year, but that may become multi-product or serve multiple markets in subsequent years. In practice, a plant must be single-product for at least two consecutive years to be included in the sample. In addition, we use selection correction in the spirit of Heckman (1979) to control for the probability of becoming a multi-product plant or serving multiple markets. We explain the details of the selection correction below. In what follows, we use only single-product plants during the estimation of the parameters of the production function.

Unobserved Product-Level Input Prices

Following de Loecker et al. (2016), we proxy for unobserved plant–product-level input prices using a function of output prices, market share, and product dummies. It is useful to break the mechanism into two components in order to understand why such a proxy works. The first component comes from the insight developed by Khandelwal (2010), which suggests that if two products in the same category have the same price, then, the product with larger market share should be of higher quality. This observation indicates that we can proxy for product quality using a function of output prices, market share, and product dummies. The second component is that higher quality products require higher quality inputs that are more expensive. Kugler and Verhoogen (2012) provide a theoretical framework for this assertion and provide supportive empirical evidence using data from Colombian manufacturing plants. Although the Mexican survey data do not offer input prices for intermediate inputs, we can test whether this result holds approximately by using data on wages. Figure 1 in Online Appendix Section 1 shows a positive correlation between average hourly wages at the plant level and output prices. Consequently, more expensive products in our sample use more costly labor inputs.

Combining these two components implies that a product with a higher market share conditional on its price should be of higher quality and therefore must be produced using more expensive inputs. We test this relationship in Figure 3, in which we examine the correlation between the residuals from a regression of market shares on output prices and product dummies (our measure of quality), and average wages.

As the figure illustrates, there is a mainly positive relationship between the residuals (the proxy for quality) and average plant-level wages. This relationship is the foundation of the input price control function that we use:

$$w_{ijt} = w_t(p_{ijt}, ms_{ijt}, D_i, G_j, EXP_{jt}; \delta_{\mathbf{s}}),$$

where p_{ijt} is the logarithm of the price of product i, ms_{ijt} is the market share of product i, D_i is a product dummy, G_j is a plant's state, EXP_{jt} is a plant's export status at time t, and δ_s is a sector-specific parameter vector that we estimate. The specification of the input price control function allows for different products manufactured by the same plant to have different input prices, and for the same product produced by different plants to have different input prices. The framework, however, does not permit separate control functions for each input. In principle, we could include an input-specific control function during estimation of the parameters of the production function, but we would be unable to identify δ_s. Unfortunately, δ_s is necessary to recover the input expenditures per product for multi-product plants.[26]

[26]See de Loecker et al. (2016) for a detailed discussion of this point.

Unobserved Plant-Level Productivity

To control for unobserved productivity, we follow the proxy methods developed by Olley and Pakes (1996) and Levinsohn and Petrin (2003) and use a control function based on a static input demand equation for materials. We use only single-product plants in the estimation, and therefore we simplify the notation at the plant level in what follows. Reference to plant j thus refers to a product in the following procedures. We assume that demand for materials takes the form:

$$\hat{m}_{jt} = m_t(\omega_{jt}, \hat{k}_{jt}, \hat{l}_{jt}, p_{jt}, ms_{jt}, D_j, G_j, EXP_{jt}, \tau_{jt}^{output}, \tau_{jt}^{input}, \tau_{jt}^{US}),$$

where $\hat{m}_{jt}$, $\hat{k}_{jt}$, and $\hat{l}_{jt}$ denote expenditures on materials, capital, and labor deflated by their respective industry price index, τ_{jt}^{output} is the tariff applied to the product produced by plant j, τ_{jt}^{input} is the tariff applied to the intermediate inputs used by plant j, and τ_{jt}^{US} is the U.S. tariff applied to the product produced by plant j. Under the assumption that demand for materials is increasing with productivity, we could potentially invert the demand function to arrive at a control function for productivity:

$$\omega_{jt} = h_t(\hat{\mathbf{x}}_{\mathbf{jt}}, \mathbf{z}_{\mathbf{jt}}),$$

with $\hat{\mathbf{x}}_{\mathbf{jt}} = (\hat{m}_{jt}, \hat{l}_{jt}, \hat{k}_{jt})$ and $\mathbf{z}_{\mathbf{jt}} = (p_{jt}, ms_{jt}, D_j, G_j, EXP_{jt}, \tau_{jt}^{output}, \tau_{jt}^{input}, \tau_{jt}^{US})$. We use second-order polynomials on $\hat{\mathbf{x}}_{\mathbf{jt}}$ and $\mathbf{z}_{\mathbf{jt}}$ to approximate the unknown function $h_t(\cdot)$ in order to control for unobserved productivity.

Selection Correction

To resolve selection bias associated with estimations using single-product-single-market producers, we use the probability of remaining as a single-product-single-market plant as a control. We assume, as in Mayer et al. (2014), that the number of products increases with productivity. Let the state vector of plant j at time t be:

$$\mathbf{s}_{jt} = (N_{jt}, K_{jt}, \Omega_{jt}, G_j, EXP_{jt}, \tau_{jt}^{output}, \tau_{jt}^{input}, \tau_{jt}^{US}),$$

where N_{jt} denotes the number of products produced by plant j at time t, and K_{jt} is the capital stock. Denote by $\bar{\omega}_{jt}(\mathbf{s}_{jt})$ the productivity cutoff associated with the introduction of a second product as a function of state variable $\mathbf{s}_{jt}$. Define the indicator variable $\mathcal{I}_{jt} = 1$ if a plant remains single-product. We can then write the probability of remaining single-product as:

$$\begin{aligned} Pr(\mathcal{I}_{jt} = 1) &= Pr(\omega_{jt} \leq \bar{\omega}_{jt}(\mathbf{s}_{jt}) | \bar{\omega}_{jt}(\mathbf{s}_{jt}), \omega_{jt-1}) \\ &= \kappa_{t-1}(\bar{\omega}_{jt}(\mathbf{s}_{jt}), \omega_{jt-1}) \\ &= \kappa_{t-1}(\hat{x}_{jt-1}, \mathbf{z}_{jt-1}) \equiv SP_{jt}, \end{aligned}$$

where the last equality comes from substituting the control function of productivity in $t-1$ and $\mathbf{z_{jt}} = (p_{jt}, ms_{jt}, D_j, G_j, EXP_{jt}, \tau_{jt}^{output}, \tau_{jt}^{input}, \tau_{jt}^{US})$. In practice, we estimate this probability using the fitted values from a probit estimation.

Estimation

We assume that productivity follows a first-order Markov process, with the law of motion:

$$\omega_{jt} = g(\omega_{jt-1,} \tau_{jt-1}^{output}, \tau_{jt-1}^{input}, \tau_{jt-1}^{US}, EXP_{jt-1}, SP_{jt,} R\&D_{jt-1}) + \xi_{jt},$$

where SP_{jt} is the fitted probability of remaining single-product, $R\&D_{jt}$ is research and development expenditures, and ξ_{jt} is the innovation to the productivity shock. The specification for the law of motion for productivity allows tariffs and export status to influence productivity but does not assume that they will necessarily affect it. The data will tell us if there is any significant correlation between productivity and these variables. We also allow research and development expenditures to affect productivity. We estimate the parameters of the production function and input price control function by constructing moments based on innovation to productivity shock ξ_{jt}. To do this, we first express ω_{jt} as a function of the data and parameters. Plugging in the input price and productivity control functions into the production function, we can write equation (5) as:

$$q_{jt} = \phi_{jt}(\hat{\mathbf{x}}_{jt}, \mathbf{z}_{jt}) + \epsilon_{jt},$$

where the function $\phi(\cdot) = f_j(\hat{\mathbf{x}}_{\mathbf{jt}}; \beta_{\mathbf{s}}) + \Lambda(w_t(p_{jt}, ms_{jt}, D_j, G_j, EXP_{jt}; \delta_{\mathbf{s}}), \hat{\mathbf{x}}_{\mathbf{jt}}; \beta_{\mathbf{s}}) + h_t(\hat{\mathbf{x}}_{\mathbf{jt}}, \mathbf{z}_{\mathbf{jt}})$ captures the output net of measurement error. Estimating this equation and recovering $\hat{q}_{jt} = \hat{\phi}_{jt}$ enables us to dispose of ϵ. In practice, we form second-order polynomials on $\hat{\mathbf{x}}_{jt}$ and $\mathbf{z}_{jt}$ to proxy $\phi(\cdot)$ and estimate the fitted values. Once we have a measure of the output net of measurement error, we can express productivity directly as a function of the data and parameters as:

$$\omega_{jt}(\beta_s, \delta_s) = \hat{\phi}_{jt} - f_j(\hat{\mathbf{x}}_{\mathbf{jt}}; \beta_{\mathbf{s}}) - \Lambda(w_t(p_{jt}, ms_{jt}, D_j, G_j, EXP_{jt}; \delta_{\mathbf{s}}), \hat{\mathbf{x}}_{\mathbf{jt}}; \beta_{\mathbf{s}}),$$

where the input price control function has been evaluated in $\Lambda(\cdot)$. We approximate $\Lambda(\cdot)$ using a second-order polynomial on the elements of the input price control function $w_t(\cdot)$ and their interactions with input expenditures.[27] Finally, we form the moment conditions using the innovation to productivity shock:

$$\xi_{jt}(\beta, \delta) = \omega_{jt}(\beta, \delta) - E[\omega_{jt}(\beta, \delta) | \omega_{ijt-1}(\beta, \delta), \tau_{jt-1}^{output}, \tau_{jt-1}^{input}, \tau_{jt-1}^{US}, EXP_{jt-1}, SP_{jt,} R\&D_{jt-1}].$$

Following Ackerberg et al. (2015), we estimate both the parameters of the production

[27]Estimating interactions between product and state dummies and input expenditures is infeasible, so they have been excluded.

function β and the input price control function δ by GMM using the moment conditions:

$$E[\xi_{\mathbf{jt}}(\beta_{\mathbf{s}}, \delta_{\mathbf{s}})\mathbf{I}_{jt}] = 0, \tag{8}$$

where the instrument matrix $\mathbf{I_{jt}}$ includes lagged materials, current capital, current labor, and their higher order interactions. It also incorporates lagged market shares, lagged tariffs, lagged prices, lagged export status, and the interaction of lagged prices with inputs and market shares. We also include a time trend and its square to control for aggregate macroeconomic trends.

Estimation yields consistent estimates of the parameters of the production function β and input price control function δ. Identification of these parameters comes from the timing assumptions on productivity. We assume that labor and capital do not respond contemporaneously to the innovation to productivity shock, but materials do, in order to construct the appropriate moment conditions. We follow de Loecker et al. (2016) and assume that input and output prices are contemporaneously correlated with innovations to productivity to construct the moments needed to identify the parameters of the input price control function.

As mentioned above, in principle, we would like to estimate the production function and input price control function at the product level, but in practice, we do not have enough observations of single-product plants that produce each product. Therefore, we follow the literature and estimate the production functions and input price control function at the sector level. We use the following sectors in the estimation: food and beverage, textiles, apparel, wood and furniture, paper industries, chemical industries, non-metallic mineral products, metallic manufacturing, and machinery and transportation equipment.

3.3 Recovering the Input Expenditure Shares by Product

The procedure described above allows for estimation of the production function and input price control function parameters. With these estimates, we have all of the elements needed to construct elasticities and markups for single-product plants. However, for multi-product plants, we still need to recover input expenditure shares by product. Given that around 73% of the plants in the sample are multi-product plants, recovering input expenditure shares by product is very important if we want to analyze the impact of NAFTA on the whole manufacturing industry. Our strategy is to use the parameter estimates and the restrictions imposed by the structural model of production to recover input expenditure shares for each product. To illustrate the method of recovering input expenditure shares, we consider a simplified example with no input price differences below.[28]

Consider a plant that manufactures three products, q_1, q_2, and q_3, with a translog production function. Let ρ_i denote the logarithm of the input expenditure share of product i (common across inputs within a product), and m_i, l_i, and k_i the logarithms of the units of materials, labor, and capital used during the production of product i. Under these assumptions, we can write the logarithm of labor used in product i as $l_i = \rho_i + l$. The same is true for the other two inputs. The production function is then:

$$\begin{aligned}
q_i &= \beta_m m_i + \beta_{mm} m_i^2 + \beta_l l_i + \beta_{ll} l_i^2 + \beta_k k_i + \beta_{kk} k_i^2 + \beta_{ml} m_i l_i + \beta_{mk} m_i k_i + \beta_{lk} l_i k_i + \beta_{mlk} m_i l_i k_i + \omega \\
&= \beta_m (m + \rho_i) + \beta_{mm} (m + \rho_i)^2 + \beta_l (l + \rho_i) + \beta_{ll} (l + \rho_i)^2 + \beta_k (k + \rho_i) + \beta_{kk} (k + \rho_i)^2 + \\
&\quad \beta_{ml} (m + \rho_i)(l + \rho_i) + \beta_{mk} (m + \rho_i)(k + \rho_i) + \beta_{lk} (l + \rho_i)(k + \rho_i) + \beta (m + \rho_i)(l + \rho_i)(k + \rho_i) \\
&\quad + \omega
\end{aligned}$$

for $i = \{1, 2, 3\}$, where m, l, and k are logarithms of the total units of materials, labor, and

[28] The intuition and procedure work for a more general production function and plant–product input price differences. See Online Appendix Section 2.

capital used by the plant. After manipulation, we have the following equations:

$$
\begin{aligned}
q_1 - \beta_m m - \beta_l l - \beta_k k - \beta_{mm} m^2 - \beta_{ll} l^2 - &= A\rho_1 + B\rho_1^2 + C\rho_1^3 + \omega, \\
\beta_{kk} k^2 - \beta_{ml} m - \beta_{mk} mk - \beta_{lk} lk - \beta_{mlk} mlk & \\
q_2 - \beta_m m - \beta_l l - \beta_k k - \beta_{mm} m^2 - \beta_{ll} l^2 - &= A\rho_2 + B\rho_2^2 + C\rho_2^3 + \omega, \\
\beta_{kk} k^2 - \beta_{ml} m - \beta_{mk} mk - \beta_{lk} lk - \beta_{mlk} mlk & \\
q_3 - \beta_m m - \beta_l l - \beta_k k - \beta_{mm} m^2 - \beta_{ll} l^2 - &= A\rho_3 + B\rho_3^2 + C\rho_3^3 + \omega, \\
\beta_{kk} k^2 - \beta_{ml} m - \beta_{mk} mk - \beta_{lk} lk - \beta_{mlk} mlk &
\end{aligned}
\tag{9}
$$

with:

$$
\begin{aligned}
A &= \beta_m + \beta_l + \beta_k + 2\beta_{mm} m + 2\beta_{ll} l + 2\beta_{kk} k + \beta_{ml} m + \beta_{ml} l + \beta_{mk} m \\
&\quad + \beta_{mk} k + \beta_{lk} l + \beta_{lk} l + \beta_{mlk} ml + \beta_{mlk} mk + \beta_{mlk} lk, \\
B &= \beta_{mm} + \beta_{ll} + \beta_{kk} + \beta_{ml} + \beta_{mk} + \beta_{lk} + \beta_{mlk} m + \beta_{mlk} l + \beta_{mlk} k, \\
C &= \beta_{mlk}.
\end{aligned}
$$

Given the estimates of the production function parameters (the βs), the left-hand side of the system of equations (9) depends only on product quantities and plant-level inputs, both of which are included in the dataset. The right-hand side depends on the unknown input expenditure shares ρ_i and unobserved productivity parameter ω. We thus have a nonlinear system of three equations and four unknowns. To solve the system, we need an additional equation, given by the assumption that all inputs used by a plant are allocated to the manufacturing of products. Under this assumption, we must have:

$$\sum_{i=1}^{3} \exp(\rho_i) = 1.$$

This final equation completes a system of four nonlinear equations and four unknowns. We can numerically solve it to recover both unobserved input expenditure shares and plant-level productivity for multi-product plants.[29] This is where the assumption of productivity being at the plant level is important. Were productivity allowed to be product-specific, the previous system would have four equations but six unknowns, making it impossible to solve.

Returning to the general framework, we can construct output elasticities for all the products in the sample by recovering the input expenditure shares for each product. We can then use these estimates to construct markups from $\hat{\mu}_{ijt} = \hat{\theta}^M_{ijt} \left(\frac{\exp(\hat{\rho}_{ijt})\hat{M}_{jt}}{P_{ijt}Q_{ijt}} \right)^{-1}$, where $\hat{\theta}^M_{ijt}$ is the estimated output elasticity of materials for product i from plant j at time t, $\hat{\rho}_{ijt}$ is the corresponding logarithm of the input expenditure share, and $\hat{M}_{jt}$ denotes the plant's expenditure on materials. We can then use data on prices to construct an estimate of marginal costs from $M\hat{C}_{ijt} = \frac{P_{ijt}}{\hat{\mu}_{ijt}}$.

4 Estimates of Elasticities and Markups

This section presents and analyzes the estimated elasticities and markups using the framework established in Section 3.

4.1 Estimates of Elasticities

Columns 1 through 4 of Table 5 show the median estimates of the output elasticities and returns to scale by sector (Table 9 in the Online Appendix shows the corresponding average elasticities and standard deviations). Column 5 indicates the number of observations used during estimation of the production function. Only products manufactured by single-product plants are used.

The interpretation of the magnitudes of the elasticities is as follows: an elasticity of materials equal to 0.86 means that a 1% increase in material inputs will increase output by 0.86%. Our estimates of the elasticities are in line with the results of other studies using

[29] We have a system of nonlinear equations for each multi-product plant and each year in our sample. We, therefore, have to solve for around 46,000 systems of nonlinear equations to recover input expenditures by product.

product-level data, with the largest elasticity for materials, followed by labor, and a small elasticity for capital.[30] From Column 4, we can see that most of the sectors are characterized by increasing returns to scale, with few exceptions.

To analyze the validity of estimates further, we examine the correlation between reported input expenditure shares out of total input expenditures at the plant level and the theoretical expenditure shares implied by the output elasticities under cost minimization. For production, $F(L, K, M)$, cost minimization yields a labor expenditure share equal to $\frac{\theta^L}{\theta^L+\theta^K+\theta^M}$, where θ^L, θ^K, and θ^M are the output elasticities of labor, capital, and materials, respectively. The expression for capital and materials share is analogous.

Figure 4 shows the relationship between the observed material expenditure share out of total input expenditure and the implied expenditure share given by $\frac{\hat{\theta}^M}{\hat{\theta}^L+\hat{\theta}^K+\hat{\theta}^M}$, where $\hat{\theta}^L$, $\hat{\theta}^K$, and $\hat{\theta}^M$ are the estimates of the output elasticities of labor, capital, and materials, respectively. As the figure shows, there is a positive and significant relationship between the observed expenditure shares of materials out of total inputs and the expenditure share implied by our estimated elasticities. We consider this as additional support for the plausibility of our elasticity estimates. Figure 4 in the Online Appendix shows a similar positive relationship for the capital and labor share.

To analyze the importance of the input price control function and the selection correction in the estimation, we re-estimated the production functions and calculated the resulting output elasticities without them. We find that correcting for input price differences during estimation is crucial. As Table 6 shows, exclusion of the input price control function yields highly implausible estimates for the elasticities, with many sectors having negative median elasticities.[31] By contrast, failure to correct for the sample selection bias that arises from conditioning on single-product plants yields estimates that are very similar to those obtained during benchmark specification, suggesting that the selection problem is not worrisome.

[30]See, for example, de Loecker et al. (2016) and Garcia-Marin and Voigtländer (2016).

[31]The same is true for average elasticities. See Table 10 in the Online Appendix.

4.2 Estimates of Markups

Table 7 shows the median estimates of markups by sector and destination.[32] A markup of 1.12 in the table means that that the price has a 12% markup above marginal cost. The value of the estimated markups is reasonable and similar to those found in studies by de Loecker et al. (2016) in India and Garcia-Marin and Voigtländer (2016) in Chile, both of which use product-level data. Our finding that the median markup on exported goods is lower than the median markup on domestically sold goods is also consistent with the findings of Blum et al. (2018) using Chilean data. In Table 7 and all of the results that follow, we have trimmed outliers above the 99th and below the 1st percentiles of the markup distribution by sector and destination to make sure that outliers are not driving the main results. All the results presented in the paper, however, are robust to using the original untrimmed sample or to trimming the top 97th and bottom 3rd percentile of the markup distribution.[33]

To verify that the estimates capture markups accurately, we compare them to accounting measures of revenue/variable costs at the plant level. Figure 5 shows a positive and significant correlation between the estimates of markups and accounting measures of revenue/variable costs, supporting the validity of the estimates. We also confirm that the estimates for markups and marginal costs are consistent with the return to scale parameters. Since most of the sectors are characterized by increasing return technologies, we expect a negative correlation between quantities produced and marginal costs. Figure 6 confirms this prediction, suggesting that plants that produce larger amounts, on average, have lower marginal costs and higher markups.

The estimates also concur with theoretical models in the multi-product plant literature, such as Mayer et al. (2014). Consistent with the predictions of their model, we find that a plant's most important products, measured by revenue share, have lower marginal costs and that plants charge a higher markup on such products (see Figure 3 in the Online Appendix).

[32] Table 11 in the Online Appendix shows the average and median markups by sector.

[33] See Section 4.1 in the Online Appendix.

Having established the validity of our estimates of elasticities, markups, and marginal costs, in the following section, we use these estimates to analyze the impact of NAFTA.

5 The Effects of Tariff Reductions

In this section, we use the product–destination-level estimates of markups and marginal costs to analyze the impact of tariff declines from NAFTA on prices through their impact on markups and marginal costs, using data for the period 1994–2008. We first analyze both the average effects of tariff reductions and these effects over time. Moreover, we also examine the effect of tariffs on estimated markups after controlling for marginal costs in order to assess whether tariff reductions affected competition.

5.1 The Impact of Tariff Reductions

5.1.1 The Average Impact of Tariff Reductions

To directly analyze the impact of tariff reductions during NAFTA on prices, markups, and marginal costs of products, we estimate the following regression equation:

$$\log Y_{ijt} = \alpha + \beta_1 \tau_{it}^{output} + \beta_2 \tau_{jt}^{input} + \beta_3 \tau_{it}^{US} + \xi_{ij} + \psi_{st} + \varepsilon_{ijt}, \tag{10}$$

where Y_{ijt} is either the price, markup, or marginal cost of product i from plant j at time t. τ_{it}^{output} is the Mexican output tariff applied to product i, τ_{jt}^{input} is a weighted measure of Mexican tariffs applied to the intermediate inputs of plant j, and τ_{it}^{US} is the tariff applied by the United States on good i coming from Mexico. ξ_{ij} and ψ_{st} are plant–product and sector–year fixed effects, respectively. We estimate equation (10) separately for domestic and exported products, and identify the coefficients on tariffs of equation (10) from variation in time of the dependent variables and tariffs within a plant–product pair.[34] We also cluster

[34]We exclude outliers in the top and bottom 1% of the markup distribution within each sector in the regressions. The results, however, are robust to the use of the full sample or the use of other cutoffs. We also

standard errors at the product level because that is the level of variation of the majority of the policy variables. The results are robust to clustering standard errors by different groups.[35]

To assess the mechanism by which tariffs affect prices, we decompose the impact of changes in tariffs on the prices of domestic products into their impact on marginal costs and markups and present the estimation results in Table 8. Column 1 repeats the estimates of the effect of tariffs on the prices of domestic products shown in Column 1 of Table 4 in Section 2.3. Columns 2 and 3 show the results of the impact of tariff changes on marginal costs and on markups, respectively. Because the logarithm of prices is exactly equal to the sum of the logarithm of marginal costs and markups, the sum of the coefficients in Columns 2 and 3 must be equal to the coefficients in Column 1. In this way, we can analyze the impact of each tariff on prices through their impact on marginal costs and markups.

We first analyze the impact of output tariffs. From Column 2, the coefficient for output tariffs on the marginal cost specification is positive and statistically significant. The magnitude of the coefficient implies that a one percentage point decline in output tariffs led to a reduction in marginal costs of 0.15%. One channel that explains this result is competition. Many studies document that increases in foreign competition, measured by cuts in output tariffs, increase plant-level productivity, as it forces plants to reduce their X-inefficiencies.[36] Using a sample of Mexican plants for 1993–2000, López-Córdova (2003) finds that a one percentage point decline in output tariffs raised revenue total factor productivity by 0.26%. In Appendix A.3, we estimate the impact of tariff declines on our measure of quantity total factor productivity (TFPQ). Consistent with the results of López-Córdova (2003), we find that a one percentage point decline in output tariffs led to a 1.3% increase in TFPQ. Since

conduct a different estimation specification from equation (10) and the main results still hold. See Section 4 in the Online Appendix for details.

[35]Since markups and marginal cost measures are estimates themselves, ideally, we would construct standard errors using bootstrapping methods to take into account potential measurement errors and parameter uncertainty in the first stage of the estimation. In particular, we would bootstrap over the entire procedure to construct standard errors of the coefficients in equation (10). However, this procedure is infeasible due to the capacity limits of INEGI's computers.

[36]See López-Córdova (2003), Pavcnik (2002), and Khandelwal and Topalova (2011).

productivity and marginal costs are negatively correlated (see Online Appendix Section 2), the rise in productivity might be driving the correlation between marginal costs and output tariffs that we find in the data. The 0.15% decline in marginal costs is not passed through completely to prices, which declined by 0.13%. This led to a small increase in the markup of 0.02%, which is not statistically significant.

The lack of statistical significance in the markups, however, does not mean that there are no pro-competitive effects. The coefficient for output tariffs in the specification for the markup captures two distinct effects. One the one hand, it captures the pro-competitive forces whereby an increase in foreign competition compelled Mexican domestic producers to lower their markups. On the other hand, it also captures the fact that declines in output tariffs led to reductions in marginal costs that were not completely passed through to prices and hence increased the markup. The fact that the coefficient is negative, although not statistically significant, suggests that the incomplete pass-through effect dominates.

The coefficients for intermediate input tariffs are not statistically significant in any of the specifications. One of the reasons for this is because the majority of plants in the sample did not use imported intermediates in their production. Only 49% of the plants in the sample used imported intermediates, and the fraction among non-exporters was even lower at 35%. Moreover, because specification (10) captures the average impact of tariff changes across time, it might not be able to capture cost savings for plants that change their intermediate input suppliers from domestic to foreign sources with a lag. Once we condition on the sample of importing plants, however, we do indeed find that declines in intermediate input tariffs led to a large and statistically significant decline in marginal costs and an increase in markups.[37]

Another explanation for the lack of statistical significance is that a decline in tariffs on intermediate inputs might have increased the incentives for local producers to upgrade the quality of their products. To the extent that quality upgrading is costly, this would be reflected in higher marginal costs, driving a negative correlation between intermediate input

[37]See Table A1 in the Appendix.

tariffs and costs. Indeed, as we show in Section 5.1.2 where we explore the time-varying effects of intermediate input tariff reductions, we find that the marginal costs gradually increased after an initial decline, and the quality estimates increased gradually over time in response to the input tariff reductions.

Finally, we analyze the impact of declines in U.S. tariffs on the three dependent variables. The coefficient in the marginal cost specification is positive and statistically significant. Its magnitude suggests that a one percentage point decline in U.S. tariffs led to a reduction in marginal costs of 1.06%. An explanation for this result is that increased market access to the U.S., captured by declines in U.S. tariffs, raised the productivity of plants. López-Córdova (2003), for example, finds that a one percentage point decline in U.S. tariffs applied to Mexican products increased revenue total factor productivity by 3%. Consistent with these results, we show in Appendix A.3 that declines in U.S. tariffs led to a rise in TFPQ of around 3.4% in our preferred specification. Increased market access might have led to an increase in productivity through two channels. First, it increased the incentives for Mexican producers to upgrade their technologies to export to the U.S. market. Second, the expansion of markets enabled producers to increase their production, allowing them to exploit economies of scale. Supporting the first channel, Iacovone and Javorcik (2012) find that Mexican plants significantly increased investment in the years immediately before the introduction of new exported products. Complementing their evidence, in Appendix A, we find that reductions in U.S. tariffs significantly increased the probability of introducing a product to the export market. In support of the second channel in Table 5 in the previous section, we find that the median product in the sample operates under increasing returns to scale. The large impact of a one percentage point reduction in U.S. tariffs on marginal costs is not passed through to prices, leading to a significant increase in the markup of around 1.01%

Using the estimated coefficients in Table 8 and the observed average decline in tariffs during this period, we calculate that tariff declines during NAFTA led to an 8.24% reduction in the marginal costs of domestic products. This decline only led to a reduction in prices of 3.43%, with markups increasing by 4.81%.

We next analyze the impact of tariff declines on exported products, summarized in Table 9. We first examine the impact of output tariffs. From Column 2, the effect of output tariffs on the marginal costs of exported products is positive and statistically significant. The reason for this result is that manufacturers of around 96% of exported products in the sample produced the same product for the domestic market. Declines in output tariffs, therefore, raised competition on their domestic products, compelling them to increase plant-level productivity and ultimately leading to a reduction in the marginal costs of their exported products. The magnitude of the coefficient suggests that a one percentage point decline in output tariffs reduced the marginal cost of exported products by 0.52%.

Looking at intermediate input tariffs, we find that reductions in input tariffs had a substantial impact on the marginal costs of exported products. The coefficient suggests that a one percentage point decline in intermediate input tariffs reduced the marginal costs of exported products by 1.45%. Declines in intermediate input tariffs can have a large impact on the marginal costs of exported products via two channels. The first channel is the direct impact through the reduction in the cost of intermediates: for firms that use imported intermediates in their production, a decline in tariffs directly reduces their costs. The second channel is through the impact of cheaper intermediates on productivity. As documented by Amiti and Konings (2007) and Khandelwal and Topalova (2011), the availability of cheaper imported inputs can raise productivity via learning, variety, and quality effects. These two channels work in the same direction leading to a significant reduction in marginal costs. Krueger (2004) provides some anecdotal evidence supporting the second channel for the largest Mexican producer of refrigerators. She documents that before NAFTA, Mexican refrigerators were of poor quality and had a very short lifespan. The reason for the short durability was the low quality of Mexican-made compressors. Once NAFTA came into effect, the manufacturer was able to substitute its Mexican compressor with a better made American compressor at competitive prices. The subsequent increase in productivity observed by the manufacturer allowed it to become the largest supplier of smaller, apartment-sized refrigerators to the U.S. market. We find some evidence of this channel in Appendix A.3, with reductions in intermediate input tariffs leading to an increase in TFPQ of around 2.2%. The contrast between domestic and exported products can be explained by the fact that the latter used more imported intermediates. Among exporters, the share of plants using imported intermediates was 72% versus 35% for non-exporters. Because of incomplete pass-through of costs to prices, a one percentage point decline in intermediate input tariffs led to a 1.46% increase in the markup.

Finally, we analyze the impact of cuts in U.S. tariffs. We find a positive and significant coefficient for U.S. tariffs on marginal costs. Its magnitude implies that a one percentage point decline in U.S. tariffs led to a 1.34% reduction in marginal costs. The intuition for the decrease in marginal costs is the same as that for domestic products above, namely that increased market access raises productivity. Looking at the impact of U.S. tariffs on the price of exported products, we find that a one percentage point decline in U.S. tariffs led to a 0.5% increase in the price of exported products. Finally, a one percentage point decline in U.S. tariffs led to a 1.84% increase in the markups. Markups increased by a significant amount because of two forces. On the one hand, reductions in marginal costs raised the markup because of incomplete pass-through. On the other hand, declines in U.S. tariff raised prices and therefore markups, as exporters could raise prices without increasing the after-tariff price paid by U.S. consumers.

The estimates suggest that the observed average decline in tariffs during this period led to a 27.93% reduction in the marginal costs of exported products, a non-significant 1.98% increase in prices, and a 29.92% increase in markups.

5.1.2 The Impact of Tariff Reductions Over Time

In the previous section, we used specification (10) and explored the average impact of changes in tariffs from NAFTA on prices, markups, and marginal costs. Since there may be a lag in plants' response to changes in tariffs, the results may not capture the full impact of tariff declines from NAFTA. This might be the case, for example, if it takes time to integrate cheaper imported intermediates into the production process, if it takes time for foreign competition to respond to changes in tariffs, or if it takes time to upgrade the quality of the product.

To shed light on these effects that may vary over time, we follow the analysis of Dix-Carneiro and Kovak (2017) and analyze the impact of tariff reductions from NAFTA through time. We consider the following specification:

$$\log Y_{ijt} - \log Y_{ij1994} = \alpha + \beta_{1t}\Delta\tau^{output}_{i,93-08} + \beta_{2t}\Delta\tau^{input}_{j,93-08} + \beta_{3t}\Delta\tau^{US}_{i,93-08} + \varphi_{st} + \epsilon_{ijt}, \quad (11)$$

for $t = \{1995, 1996, ..., 2008\}$. The term $\log Y_{ijt} - \log Y_{ij94}$ is the log change in prices, markups, and marginal costs between year t and 1994, and $\Delta\tau_{93-08}$ is the 93–08 change in tariff measures. φ_{st} is the sector fixed effects, and the coefficients β_{1t}, β_{2t}, and β_{3t} capture the cumulative effect of tariff declines by year t.

Figure 7 shows the evolution of the effects of output tariffs on prices, marginal costs, and markups for both domestic and exported products. Consistent with the results for the average effects, the price of domestic goods declined in response to the reduction in output tariffs, and this decline was mostly driven by the fall in marginal costs. However, the results indicate that while the drop in marginal costs due to the decline in output tariffs materialized immediately, prices fell only gradually. For exported products, prices were relatively stable over time and marginal costs decreased in a more gradual manner.

We then show the evolution of the effects of intermediate input tariffs on the three measures for domestic and exported products in Figure 8. In contrast with the results on the average effects, we see a clear dynamic pattern in the effects of intermediate input tariffs on domestic products. Initially, the reduction in intermediate input tariffs leads to a decline in prices. This decline is driven by the reduction in marginal costs, from the fall in the price of intermediate inputs. Over time, however, reductions in intermediate input tariffs appear to lead to an increase in the prices of domestic products. For exported products, we observe similar dynamic patterns in the effect of input tariffs, but changes are again more gradual.

A possible explanation for the patterns in Figure 8 is that over time, Mexican firms chose to upgrade the quality of their products by using higher quality intermediates. This is also consistent with the pattern of marginal costs, which increases over time and partially offsets the initial decline in the prices of intermediate goods. Markups gradually increase over time, and this is consistent with Mexican firms upgrading the quality of their products as long as the higher quality products have higher markups. To verify this, we take the residuals from a regression of market shares on output prices and product dummies and consider their evolution as the left-hand side of specification (11). This proxy of quality follows the intuition introduced in Khandelwal (2010): conditional on prices, products with higher market shares are assigned better quality.[38] Figure 9 shows that the quality of domestic products indeed increased gradually in response to the decline in input tariffs. On the other hand, the quality improvement of exported products was weaker and mostly insignificant over time.

Finally, we turn to Figure 10, in which we present the evolution of the effects of U.S. tariffs on the three measures. Consistent with the results on the average effects, we see an increase in markups and a decline in the marginal costs of domestic products in response to the decline in U.S. tariffs. A similar pattern is also found in the markups of exported products. However, marginal costs did not decline as significantly, resulting in an increase in prices. In addition, all these movements did not happen immediately and took around a decade to materialize.

[38] In this case, identifying the dynamic effect of tariffs on quality requires the unobserved product-time specific component of quality to be orthogonal to tariffs. More details on our quality estimation and dynamic changes in quality as a response to other tariff changes are reported in Online Appendix Section 5.

5.2 Effects of Competition on Markups

The point estimates presented in Column 3 of Table 8 seem to suggest that declines in output tariffs did not have an impact on the markup of domestic products. Considering only these point estimates as evidence of pro-competitive effects, however, is misleading, since the specification in Column 3 does not hold marginal costs fixed and therefore might not be capturing the *ceteris paribus* impact of a decline in output tariffs on markups. If plants adjusted their markups in response to changes in marginal costs, the coefficient in Column 3 would capture not only the direct impact of tariffs on the markups, but also the indirect impact through changes in marginal costs that are driven by changes in tariffs. In fact, the nature of variable markups implies that the covariance between markups and marginal costs is non-zero. Only in the particular case of complete pass-through of marginal costs to prices would this covariance be zero.[39] The finding that there is a positive correlation between marginal costs and output tariffs suggests that the coefficient in Column 3 indeed captures these two effects. Intuitively, changes in output tariffs lead to adjustments in marginal costs that are not entirely passed through to prices, causing movements in markups. The coefficient in Column 3, therefore, captures both the direct impact of output tariffs on markups (the pro-competitive effects) and the indirect impact of marginal costs on markups.

To test for effects of competition on markups, we control for marginal costs. We estimate the following regression:

$$\log \mu_{ijt} = \alpha + \kappa_1 \tau_{it}^{output} + \kappa_2 \tau_{it}^{US} + g(\hat{mc}_{ijt}; \eta) + \xi_{ij} + \psi_{st} + \varepsilon_{ijt}, \tag{12}$$

where $g(\hat{mc}_{ijt}; \eta)$ is a polynomial on marginal costs used as a control. The specification also includes plant–product fixed effects, ξ_{ij}, and sector–year fixed effects, ψ_{st}. We use a third-order polynomial on marginal costs, but the results are robust to a linear or second-order polynomial. Measurement errors in marginal costs will lead to attenuation bias. We,

[39] See Online Appendix Section 2 for the formal argument. In Table 8 in the Online Appendix, we provide evidence of incomplete pass-through of marginal costs to prices.

therefore, instrument the polynomial of marginal costs with its lagged polynomial and intermediate input tariffs. These instruments are valid as long as there is no serial correlation in the measurement errors of marginal costs. We estimate equation (12) separately for domestic and exported products.[40]

Table 10 shows results from the estimation of equation (12) for the sample of domestic products. We present the baseline specification in Column 1 and instrument for the polynomial on marginal costs using the lagged polynomial and input tariffs in Column 2. From Column 1, once we control for marginal costs, the coefficient for the impact of output tariffs on markups becomes positive and statistically significant. Its magnitude suggests that a one percentage point decrease in output tariffs led to a 0.11% decline in the markup of domestic products. Thus, output tariff declines during this period had pro-competitive effects. The impact of U.S. tariffs on domestic products is, on the other hand, insignificant. Contrasting this result with the one from Column 3 of Table 8, we can see that all of the impact of U.S. tariffs on the markups of domestic products comes through their impact on marginal costs. Once we control for marginal costs, U.S. tariffs have no additional effect on the markup of domestic products. From Column 2, instrumenting the marginal-cost polynomial affects the values of the coefficients only slightly, leaving conclusions above unaltered.

For exported products, the coefficient for output tariffs in Column 1 of Table 11 is not statistically significant. This result shows that the impact of output tariffs on the markup of exported products found in Table 9 comes from its effect on marginal costs. Holding marginal costs fixed, output tariffs have no additional effect on the markup of exported products. We obtain a negative and significant effect of U.S. tariffs on the markups of exported products. A one percentage point decline in U.S. tariffs led to a 0.54% increase in the markup of exported products. This last result is consistent with the anti-competitive effects of trade liberalization predicted by the model constructed in de Blas and Russ (2015). This work is the first to document this mechanism empirically.

[40] In Appendix A.1, we explore the heterogeneity of results along a number of dimensions.

5.3 Discussion

The results in the previous section suggest that tariff declines from NAFTA led both to an increase in competition in the domestic market, forcing domestic producers to lower their markups pushing down prices, and separately, to a reduction in marginal costs that was not completely passed through to prices. The reason for this incomplete pass-through is that producers offset part of the decline in marginal costs by raising their markups. Using the coefficients from Column 3 in Table 8 and the observed reduction in tariffs suggests that markups increased by 4.81% during this period. The fact that we observed a rise in markups, however, does not mean that competitive forces were not significant. Without competitive forces, it is likely that the prices of domestic products would have declined by less as producers captured more of the gains from declines in marginal costs. The estimates suggest that while Mexican consumers benefited from a decrease in prices, producers profited by more, as the reduction in marginal costs, induced by reductions in tariffs, allowed them to increase their margins despite the increased competitive pressure. Finally, note that while the total declines in prices were small relative to the size of the tariff reductions, we cannot rule out quality upgrading as another potential source of consumer gains.

For exported products, the results suggest that tariff reductions during NAFTA led to an increase in markups as cuts to U.S. tariffs on Mexican products allowed Mexican producers to raise their prices and because reductions in marginal costs were not passed through to prices as producers offset the decline in costs by raising markups. These two effects led to substantial improvements in the profitability of Mexican exporters. The fact that producers increased their markups in response to declines in U.S. tariffs, however, does not mean that U.S. consumers did not benefit from NAFTA. First, reductions in U.S. tariffs led to an increase in the prices of exported products that was smaller than the actual decline in tariffs. This suggests that after-tariff prices paid by U.S. consumers still declined. Second, the period was characterized by introductions of new products to the export market, with approximately 15% new exported products introduced each year on average (see Table 6 in the Online Appendix). Therefore, the U.S. consumers significantly benefited from the extensive margin through the availability of new product varieties.

Moreover, we also explore the heterogeneity of our main results and the detailed analysis is presented in Appendix A.1. In particular, we explore how these pro-competitive and anti-competitive effects vary across several dimensions such as type of goods and firm size. We find that pro-competitive effects were stronger for homogeneous products (using the classification in Rauch (1999)), whereas anti-competitive effects were mostly associated with non-homogeneous exported goods. In addition, pro-competitive effects tended to come from large firms, but anti-competitive effects were more significant among small exporters.

All in all, the results suggest that Mexican manufacturers were the definite beneficiaries of NAFTA. They were able to increase their profitability despite the increase in foreign competition which allowed them to boost their margins. Moreover, it appears that exporters disproportionately benefited, as the reductions in U.S. tariffs on Mexican products permitted an additional expansion in their profit margins.

6 Conclusion

This paper analyzed the impact of NAFTA on prices and competition in Mexico using detailed disaggregated information on Mexican manufacturing plants. Employing the methodology developed by de Loecker et al. (2016), we estimated markups and marginal costs at the plant–product–destination level by exploiting quantity and price data on domestic and exported products. Our results suggest that tariff declines from NAFTA led to large reductions in the marginal costs of domestic and exported products through different channels. Declines in output tariffs lowered the marginal costs of plants as the increase in foreign competition compelled them to raise their productivity. Cuts in tariffs on intermediate inputs directly reduced the marginal costs of plants that used imported intermediates. Finally, declines in U.S. tariffs lowered marginal costs as the increased market access allowed plants to expand production and led them to upgrade their technology. These large reductions in marginal costs, however, were not passed through to lower prices, as exporters significantly increased their markups in response to improved market access.

Controlling for changes in marginal costs, the results suggest that declines in output tariffs had the expected pro-competitive effects: An increase in foreign competition compelled Mexican manufacturers to reduce markups on domestic products. We also found that exporters responded to declines in U.S. tariffs by increasing markups on exported goods. This novel evidence on anti-competitive effects through the market access channel suggests that the direct gain for consumers from a decrease in tariffs on incumbent imported goods might be lower than previously thought, as producers adjust their prices to increase their profitability.

Our analysis sheds light on how reciprocal trade liberalizations affect firms' prices and the competition firms face. The NAFTA episode and the detailed plant–product-level data from Mexican manufacturing plants provide us with an opportunity to study this question empirically. Our work complements existing literature by differentiating firms' adjustment in domestic and foreign markets. We believe that understanding firms' responses under such circumstances is very important, since a large proportion of past tariff reductions have been reciprocal.

References

Ackerberg, D. A., Caves, K., and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451.

Amiti, M. and Konings, J. (2007). Trade liberalization, intermediate inputs, and productivity: Evidence from indonesia. *American Economic Review*, 97(5):1611–1638.

Arkolakis, C., Costinot, A., Donaldson, D., and Rodríguez-Clare, A. (2017). The elusive pro-competitive effects of trade. *Working paper*.

Bernard, A. B., Eaton, J., Jensen, J. B., and Kortum, S. (2003). Plants and productivity in international trade. *American Economic Review*, 93(4):1268–1290.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

Besedes, T., Kohl, T., and Lake, J. (2018). Phase Out Tariffs, Phase In Trade?

Blum, B. S., Claro, S., Rivers, D. A., and Horstmann, I. (2018). The ABCs of Firm Heterogeneity: The Effects of Demand and Cost Differences on Exporting.

Brandt, L., Van Biesebroeck, J., Wang, L., and Zhang, Y. (2017). Wto accession and performance of chinese manufacturing firms. *American Economic Review*, 107(9):2784–2820.

Bustos, P. (2011). Trade liberalization, exports, and technology upgrading: Evidence on the impact of mercosur on argentinian firms. *American Economic Review*, 101(1):304–40.

Cajal-Grossi, J., Macchiavello, R., and Noguera, G. (2019). International Buyers' Sourcing and Suppliers' Markups in Bangladeshi Garments. CEPR Discussion Papers 13482, C.E.P.R. Discussion Papers.

Cameron, M. A. and Tomlin, B. W. (2002). *The Making of NAFTA: How the Deal Was Done*. Cornell University Press.

Caselli, M., Chatterjee, A., and Woodland, A. (2017). Multi-product exporters, variable markups and exchange rate fluctuations. *Canadian Journal of Economics/Revue canadienne d'économique*, 50(4):1130–1160.

Conconi, P., García-Santana, M., Puccio, L., and Venturini, R. (2018). From final goods to inputs: The protectionist effect of rules of origin. *American Economic Review*, 108(8):2335–2365.

de Blas, B. and Russ, K. N. (2015). Understanding markups in the open economy. *American Economic Journal: Macroeconomics*, 7(2):157–180.

de Hoyos, R. and Iacovone, L. (2013). Economic performance under nafta: A firm-level analysis of the trade-productivity linkages. *World Development*, 44:180–193.

de Loecker, J. (2011). Product differentiation, multiproduct firms, and estimating the impact of trade lberalization on productivity. *Econometrica*, 79(5):1407–1451.

de Loecker, J. and Goldberg, P. K. (2014). Firm performance in a global market. *Annual Review of Eocnomics*, 6:201–227.

de Loecker, J., Goldberg, P. K., Khandelwal, A., and Pavcnik, N. (2016). Prices, markups, and trade reform. *Econometrica*, 84(2):445–510.

de Loecker, J. and Warzynski, F. (2012). Markups and firm-level export status. *American Economic Review*, 102(6):2437–2471.

Dix-Carneiro, R. and Kovak, B. K. (2017). Trade liberalization and regional dynamics. *American Economic Review*, 107(10):2908–2946.

Easterly, W., Fiess, N., and Lederman, D. (2003). Nafta and convergence in north america: High expectations, big events, little time. *LACEA Economía*, 4(1):1–53.

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

Esquivel, G. and Rodríguez-López, J. A. (2003). Technology, trade, and wage inequality in mexico before and after nafta. *Journal of Development Economics*, 72(2):543–565.

Faber, B. (2014). Trade Liberalization, the Price of Quality, and Inequality: Evidence from Mexican Store Prices. (October).

Garcia-Marin, A. and Voigtländer, N. (2016). Exporting and plant-level efficiency gains: Its in the measure. *UCLA mimeo*.

Goldberg, P. K. (1995). Product differentiation and oligopoly in international market: The case of the u.s. *Econometrica*, 63(4):891–951.

Goldberg, P. K. and Hellerstein, R. (2013). A structural approach to indentifying the sources of local currency price stability. *Review of Economic Studies*, 80:175–210.

Hall, R. E. (1986). Market structure and macroeconomic fluctuations. *Brookings papers on Eco nomic Activity*, 17(2):285–322.

Hallak, J. C. (2006). Product quality and the direction of trade. *Journal of International Economics*, 68(1):238–265.

Hallak, J. C. and Sivadasan, J. (2013). Product and process productivity: Implications for quality choice and conditional exporter premia. *Journal of International Economics*, 9(1):53–67.

Harrison, A. E. (1994). Productivity, imperfect competition and trade reform. *Journal of International Economics*, 36:57–73.

Head, K. and Ries, J. (1999). Rationalization effects of tariff reductions. *Journal of International Economics*, 47:295–320.

Heckman, J. J. (1979). Sample selection as a specification error. *Econometrica*, 47(1):153–161.

Iacovone, L. (2012). The better you are the stronger it makes you: Evidence on the asymmetric impact of liberalization. *Journal of Development Economics*, 99(2):474–485.

Iacovone, L. and Javorcik, B. S. (2012). Getting ready: Preparation for exporting. *CEPR Discussion Paper*, (8926).

Juarez, L. and de la Cabada, D. C. (2016). Downward wage rigidities in the mexican labor market 1996-2011. *Banco de México Working Papers*, (2016-23).

Khandelwal, A. (2010). The long and short (of) quality ladders. *Review of Economic Studies*, 77:1450–1476.

Khandelwal, A. and Topalova, P. (2011). Trade liberalization and firm productivity: The case of india. *The Review of Economic and Statistics*, 93(3):995–1009.

Kowalczyk, C. and Davis, D. R. (1998). *The Regionalization of the World Economy*, chapter Tariff Phase-Outs: Theory and Evidence from GATT and NAFTA, pages 227–258. University of Chicago Press.

Krishna, P. and Mitra, D. (1998). Trade liberalization, market discipline and productivity growth: new evidence from india. *Journal of Development Economics*, 56:447–462.

Krueger, A. O. (2004). Expanding trade and unleashing growth: The prospects for lasting poverty reduction. In *IMF Seminar on Trade and Regional Integration, Dakar, Senegal.*

Kugler, M. and Verhoogen, E. A. (2012). Prices, plant size, and product quality. *Review of Economic Studies*, 79(1):307–339.

Levinsohn, J. (1993). Testing the iimport-as-market-discipline hypothesis. *Journal of International Economics*, 35:1–22.

Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *The Review of Eocnomic Studies*, 70:317–341.

Lileeva, A. and Trefler, D. (2010). Improved access to foreign markets raises plant-level productivity... for some plants. *Quarterly Journal of Economics*, 125(3):1051–1099.

Linder, S. B. (1961). *An Essay on Trade and Transformation.* Stockholm: Almqvist & Wiksell.

López-Córdova, E. (2003). Nafta and manufacturing productivity in mexico. *Economía*, 4(1):55–98.

Mayer, T., Melitz, M. J., and Ottaviano, G. I. P. (2014). Market size, competition, and the product mix of exporters. *American Economic Review*, 104(2):495–536.

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

Olley, S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297.

Pavcnik, N. (2002). Trade liberalization, exit, and productivity improvements: Evidence from chilean plants. *Review of Economic Studies*, 69:245–276.

Pierce, J. R. (2011). Plant-level responses to antidumping duties: Evidence from U.S. manufacturers. *Journal of International Economics*, 85(2):222–233.

Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of International Economics*, 48(1):7–35.

Romalis, J. (2007). Naftas and cusftas impact on international trade. *The Review of Economic and Statistics*, 89(3):416–435.

Ten-Kate, A. (1992). Trade liberalization and economic stabilization in mexico: Lessons of experience. *World Development*, 20(5):659–672.

Trefler, D. (2004). The long and short of the canada-u.s. free trade agreement. *American Economic Review*, 94(4):870–895.

Utar, H. and Ruiz, L. B. (2013). International competition and industrial evolution: Evidence from the impact of Chinese competition on Mexican maquiladoras. *Journal of Development Economics*, 105:267–287.

Verhoogen, E. A. (2008). Trade, quality upgrading, and wage inequality in the mexican manufacturing sector. *Quarterly Journal of Economics*, 123(2):489–530.

A Robustness and Additional Results on the Impact of Tariffs

This section presents results from the estimation of various specifications to test the robustness of the main empirical specification. The section also includes additional empirical results on the impact of tariffs on a number of variables.

A.1 Heterogeneity of Results

A.1.1 Main Specification

In this subsection, we analyze the heterogeneity of the results from our main specification. We focus on the heterogeneity of the results by import and export status. In particular, we estimate the following two variations of the main specification:

$$\log Y_{ijt} = \alpha + \beta_1 \tau_{it}^{output} + \beta_2 \tau_{jt}^{input} + \beta_3 \tau_{it}^{US} + Import \times [\gamma_1 \tau_{it}^{output} + \gamma_2 \tau_{jt}^{input} + \gamma_3 \tau_{it}^{US}] + \xi_{ij} + \psi_{st} + \varepsilon_{ijt} \quad (13)$$

$$\log Y_{ijt} = \alpha + \beta_1 \tau_{it}^{output} + \beta_2 \tau_{jt}^{input} + \beta_3 \tau_{it}^{US} + Export \times [\gamma_1 \tau_{it}^{output} + \gamma_2 \tau_{jt}^{input} + \gamma_3 \tau_{it}^{US}] + \xi_{ij} + \psi_{st} + \varepsilon_{ijt}, \quad (14)$$

where $Import$ is an indicator variable equal to one if the plant imports intermediate inputs and $Export$ is an indicator function equal to one if the plant exports. Tables A1 and A2 show the results for these specifications.

A.1.2 Markup Specification

In this subsection, we explore the heterogeneity of our markup specification. First, Table A3 and Table A4 show the heterogeneity of the results by separating the sample of products into homogeneous products and non-homogeneous products. The classification of products follows the one proposed inRauch (1999). Under this classification, a product is classified as homogeneous if it is traded in an organized exchange.

We also analyze the heterogeneity of the results by size. We repeat the markup specification (12), but separate the estimation sample into large plants, which are classified as those having more than 150 employees, and small plants having fewer than 150 employees. Table A5 and Table A6 show the results from these specifications.

We also analyze the heterogeneous impact of tariffs on markups by initial level of the markup. To do this, we estimate the following specification:

$$\begin{aligned}\log \mu_{ijt} =& \alpha + \beta_1 \tau_{it}^{output} + \beta_2 \tau_{it}^{US} + High\_Tariff_{0,ijt} \times [\gamma_1 \tau_{it}^{output} + \gamma_2 \tau_{it}^{US}] + g(\log MC_{ijt}) \\ &+ \xi_{ij} + \psi_{st} + \varepsilon_{ijt}, \end{aligned} \tag{15}$$

where $High\_Tariff_{0,ijt}$ is an indicator function equal to one if product $i's$ markup in 1994 was in the the top 10% of that product category. Table A7 and Table A8 show the results from the estimation.

A.2 The Impact of Tariffs on the Introduction of New Products

In this subsection, we analyze the impact of tariff changes from NAFTA on product introduction. In particular, we estimate the following probit model at the plant level:

$$Pr(Intro_{jt} = 1) = \Phi(\alpha + \beta_1 \tau_{jt}^{output} + \beta_2 \tau_{jt}^{input} + \beta_3 \tau_{jt}^{US} + EXP_{jt} \times [\beta_4 \tau_{jt}^{output} + \beta_5 \tau_{jt}^{input} + \beta_6 \tau_{jt}^{US}] + \gamma X_{jt}) \tag{16}$$

where $Intro_{jt}$ is an indicator function that is equal to 1 if plant j introduces a product in year t, and $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. τ_{jt}^{output}, τ_{jt}^{input}, and τ_{jt}^{US} are plant-level output, intermediate input, and U.S. tariffs, respectively. Output and U.S. plant-level tariffs are constructed as a sales-weighted average of the tariffs on products sold by the plant. EXP_{jt} is an indicator function that is equal to 1 if plant j is an exporter plant. Finally, X_{jt} is a vector of plant-level controls that include the plant's level of labor, capital, and material expenditures, a plant's import and export status, as well as year, state, and sector fixed effects. The tariffs are interacted with the export status indicator to capture the heterogeneous impact of tariff changes for plants that export products to the U.S.

Equation (16) is estimated using contemporaneous as well as lagged independent variables. Table A9 shows the results of the estimation. As we can see from Table A9, declines in U.S. tariffs significantly increase the probability of exporters introducing a new product. This result suggests that exporters responded to the increase in market access by introducing new products to the export markets. This result is consistent with Iacovone and Javorcik (2012), who find that a decline in U.S. tariffs stimulated investment by Mexican plants, leading them to introduce new products into the export market. Declines in input tariffs, on the other hand, decrease the probability of new product introduction. As we found in Section 5.1.2, declines in input tariffs induced manufacturing plants to upgrade the quality of products they produce. The results suggest that plants concentrated more resources on quality upgrading and less on introducing new products. The results are robust to the exclusion of plant-level controls.

A.3 The Impact of Tariffs on Quantity Total Factor Productivity

In this subsection, we study the impact of changes in tariffs arising from NAFTA on quantity total factor productivity (TFPQ). As shown in Subsection 3.3, the empirical framework used to estimate production functions and construct product-level markups allows us to recover a measure of TFPQ. TFPQ represents a more appropriate measure of productivity than the more commonly used revenue total factor productivity (TFPR), as it does not confound changes in productivity with movements in prices or markups.[41] We estimate the following equation based on the specification used by López-Córdova (2003):

[41] The majority of papers analyzing the impact of trade policy on productivity use a measure of revenue productivity. Examples include the studies by Pavcnik (2002), López-Córdova (2003), Amiti and Konings (2007), Khandelwal and Topalova (2011), and Iacovone (2012).

$$\omega_{jt} = \alpha + \beta_1 \tau_{jt}^{output} + \beta_2 \tau_{jt}^{input} + \beta_3 \tau_{jt}^{US} + \gamma X_{jt} + \epsilon_{jt} \tag{17}$$

where ω_{jt} is the logarithm of TFPQ of plant j. τ_{jt}^{output}, τ_{jt}^{input}, and τ_{jt}^{US} are plant-level output, intermediate input, and U.S. tariffs, respectively. Output and U.S. plant-level tariffs are constructed as a sales-weighted average of the tariffs on products sold by plant j. X_{jt} is a vector of plant-level controls that include the plant's import and export status, the total industry sales of plant j excluding its sales, and the Herfindahl–Hirschman Index, HHI, of market concentration in the industry of plant j, as well as year, state, and sector fixed effects. Table A10 shows the results of the estimation with different fixed effects.

As we can see from Table A10, declines in tariffs lead to large and statistically significant increases in plant-level TFPQ. The magnitudes of the coefficients in Column 3, for example, suggest that a one percentage point decline in output tariffs raises TFPQ by 1.3%, a one percentage point reduction in input tariffs increases TFPQ by 2.2%, and finally cuts in U.S. tariffs of one percentage point raise TFPQ by 3.4%.

Table 1: Average Number of Plant–Product Pair per Sector

Sector	# of Products			Avg. # of Products Per Plant		
	Total	Domestic	Exported	Total	Domestic	Exported†
	(1)	(2)	(3)	(4)	(5)	(6)
Food and Beverage	2,963	2,622	340	3.66	3.23	1.55
Textile Manufacturing	548	417	131	2.36	1.78	1.13
Apparel Manufacturing	1,240	1,091	149	3.00	2.64	1.13
Wood and Furniture	610	547	63	4.13	3.70	1.57
Paper Industries	850	752	98	2.62	2.33	1.16
Chemical Industries	2,908	2,348	561	4.11	3.31	1.66
Non-Metallic Minerals	839	708	130	2.86	2.40	1.67
Metallic Manufacturing	1,105	809	296	3.17	2.30	1.66
Machinery & Equipment	890	666	225	2.76	2.06	1.17

Notes: Columns 1–3 show the average number of plant–product pairs per sector for all years in the sample. Columns 4–6 show the average number of products per plant for all years in the sample.
†Average number of exported products for exporter plants.

Table 2: Summary Statistics on Plants

Sector	# of Plants			Average (Thousands of Dollars)				
	Total	Exporter	Single	Employees†	Sales	VA/Employees	Materials	Capital
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Food and Beverage	814	224	151	378.9	35,823.9	164.7	16,230.4	7,363.1
Textile Manufacturing	235	118	85	248.9	13,339.1	56.3	5,568.6	4,116.2
Apparel Manufacturing	410	135	143	148.9	4,864.8	25.5	2,217.0	771.4
Wood and Furniture	150	42	43	133.6	4,681.2	15.2	2,457.6	1,136.5
Paper Industries	318	81	165	210.5	16,761.4	45.5	6,907.1	5,543.5
Chemical Industries	707	339	167	244.0	24,221.0	48.6	10,123.2	6,940.8
Non-Metallic Minerals	292	79	125	212.1	19,815.1	67.6	2,938.7	9,455.3
Metallic Manufacturing	354	183	123	242.7	32,703.9	36.6	18,443.7	8,752.3
Machinery & Equipment	325	192	96	467.0	78,896.3	26.3	44,954.4	15,835.7

Notes: Includes sector-level averages of plant-level variables across all years in the sample. Units in columns 5–8 are in 1994 U.S. dollars, converted from Mexican pesos using the average 1994 exchange rate. †Number of employees.

Table 3: Summary Statistics on Tariffs

	1993			2008			Difference 2008–1993	
	Mean (1)	Median (2)	S.D. (3)	Mean (4)	Median (5)	S.D. (6)	Mean (7)	Median (8)
Output Tariffs	14.8	15.0	4.3	0.2	0.0	1.0	-14.6	−15.0
Intermediate Input Tariffs	9.4	9.2	2.3	0.1	0.0	0.0	−9.3	−9.2
U.S. Tariffs	5.2	4.3	4.5	0.1	0.0	0.5	−5.1	−4.3

Notes: The table shows the mean, median, and standard deviation of the tariffs during 1993 and 2008 in percentage points. Columns 7 and 8 report the difference in means and medians, respectively, between 2008 and 1993.

Table 4: Impact of Tariffs on Prices

	Domestic (1)	Exported (2)
τ_{it}^{output}	0.13^a	0.04
	(.04)	(.08)
τ_{jt}^{input}	0.14	−0.00
	(.54)	(.14)
τ_{it}^{US}	0.05	-0.50^c
	(.14)	(0.30)
Total Impact (%)	-3.43^a	1.98
	(1.43)	(2.50)
Within R^2	0.073	0.035
N	$145,887$	$28,738$

Notes: The dependent variable is the log of prices. Column 1 uses the sample of domestic products and Column 2 the sample of exported products. The regressions exclude outliers in the top and bottom 1% of the markup distribution within each sector. Regressions include plant–product and sector–year fixed effects. Standard errors are clustered at the product level. We calculate the total impact of tariff declines by taking the average percentage point decrease in tariffs for the 1993–2008 period, -14.6, -9.3, and -5.1 for output, input, and U.S. tariffs respectively, and multiplying them by their associated coefficients.
Significance: a (1%), b (5%), and c (10%).

Table 5: Median Elasticities by Sector, All Products

Sector	Materials (1)	Capital (2)	Labor (3)	RTS (4)	Obs. in Estimation (5)
Food and Beverage	0.86	0.06	0.19	1.12	1,781
Textile Manufacturing	0.62	0.02	0.28	0.97	992
Apparel Manufacturing	0.86	0.05	0.11	1.00	1,691
Wood and Furniture Industries	0.73	0.12	0.08	0.91	490
Paper Industries	0.95	0.01	0.30	1.21	1,968
Chemical Industries	0.65	0.09	0.27	1.03	1,995
Non-Metallic Mineral Products	0.58	0.08	0.42	1.05	1,519
Metallic Manufacturing	0.67	0.18	0.23	1.09	1,493
Machinery and Equipment Manufacturing	0.74	0.02	0.10	0.80	1,073
Total	0.76	0.08	0.18	1.05	1,073

Notes: Estimates of output elasticities of the production function for all products, domestic and exported, and all years in the sample (1994–2008). Columns 1–3 report median elasticities for each sector. Column 4 reports the median returns to scale (RTS), which is the sum of labor, capital, and materials elasticities. Column 5 reports the total number of observations used during estimation of the production function for each sector. The total corresponds to the median observation across all products and years.

Table 6: Median Elasticities by Sector, No Correction

	No Selection Correction			No Input Price Correction		
Sector	Materials (1)	Capital (2)	Labor (3)	Materials (4)	Capital (5)	Labor (6)
Food and Beverage	0.88	0.03	0.16	0.69	-0.12	0.52
Textile Manufacturing	0.69	-0.01	0.21	0.56	0.12	0.28
Apparel Manufacturing	0.87	0.05	0.10	0.52	-0.01	0.12
Wood and Furniture Industries	0.68	0.08	0.07	-0.26	-0.01	-0.39
Paper Industries	0.96	0.01	0.31	0.23	-0.30	1.04
Chemical Industries	0.64	0.11	0.27	0.35	0.38	0.81
Non-Metallic Mineral Products	0.54	0.09	0.43	0.50	0.04	0.48
Metallic Manufacturing	0.68	0.15	0.22	0.80	0.40	0.11
Machinery & Equipment	0.59	0.08	0.16	0.80	-0.64	-0.30

Notes: Estimates of the output elasticities of the production function for all products, domestic and exported, and all years in the sample (1994–2008). Columns 1–3 report median elasticities in each sector for estimation, without including selection correction for the probability of remaining single-product. Columns 4–6 report median elasticities for estimation, without including the input price control function.

Table 7: Median Markup by Destination

	Markup	
Sector	Domestic	Exported
	(1)	(2)
Food and Beverage	1.12	1.11
Textile Manufacturing	1.11	1.32
Apparel Manufacturing	1.27	1.19
Wood and Furniture Industries	1.06	0.90
Paper Industries	1.47	1.93
Chemical Industries	1.29	1.03
Non-Metallic Mineral Products	1.96	1.79
Metallic Manufacturing	1.14	1.03
Machinery and Equipment Manufacturing	1.27	1.45
Total	1.24	1.19

Notes: We trim outliers above the 99th and below the 1st percentiles of the markup distribution by sector and destination. The total corresponds to the median markup across all products and years in each sample.

Table 8: Impact of Tariffs on Prices, Marginal Costs, and Markups, Domestic Products

	$\log P_{ijt}$	$\log MC_{ijt}$	$\log \mu_{ijt}$
	(1)	(2)	(3)
τ_{it}^{output}	0.13^{a}	0.15^{b}	-0.02
	(.04)	(.08)	(.09)
τ_{jt}^{input}	0.14	0.07	0.07
	(.54)	(.16)	(.16)
τ_{it}^{US}	0.05	1.06^{a}	-1.01^{a}
	(.14)	(.32)	(.32)
Total Impact (%)	-3.43^{a}	-8.24^{a}	4.81^{b}
	(1.43)	(2.67)	(2.48)
Within R^2	0.073	0.01	0.02
N	145,887	145,887	145,887

Notes: Dependent variables in Columns 1–3 are the logs of prices, marginal costs, and markups, respectively. The regressions exclude outliers in the top and bottom 1% of the markup distribution within each sector. Regressions include plant–product and sector–year fixed effects using data for the sample of domestic products and all years (1994–2008). Standard errors are clustered at the product level. We calculate the total impact of tariff declines by taking the average percentage point decrease in tariffs for the 1993–2008 period, -14.6, -9.3, and -5.1 for output, input and U.S. tariffs respectively, and multiplying them by their associated coefficients.
Significance: a (1%), b (5%), and c (10%).

Table 9: Impact of Tariffs on Prices, Marginal Costs, and Markups, Exported Products

	$\log P_{ijt}$	$\log MC_{ijt}$	$\log \mu_{ijt}$
	(1)	(2)	(3)
τ_{it}^{output}	0.04	0.52^c	-0.48^c
	(.08)	(.29)	(.28)
τ_{jt}^{input}	-0.00	1.45^a	-1.46^a
	(.14)	(.33)	(.35)
τ_{it}^{US}	-0.50^c	1.34^b	-1.84^a
	(0.30)	(.67)	(.66)
Total Impact (%)	1.98	-27.93^a	29.92^a
	(2.50)	(6.40)	(6.34)
Within R^2	0.035	0.034	0.032
N	28,738	28,738	28,738

Notes: Dependent variables in Columns 1–3 are the logs of prices, marginal costs, and markups, respectively. The regressions exclude outliers in the top and bottom 1% of the markup distribution within each sector. Regressions include plant–product and sector–year fixed effects using data for the sample of exported products and all years (1994–2008). Standard errors are clustered at the product level. We calculate the total impact of tariff declines by taking the average percentage point decrease in tariffs for the 1993–2008 period, -14.6, -9.3, and -5.1 for output, input, and U.S. tariffs respectively, and multiplying them by their associated coefficients.
Significance: a (1%), b (5%), and c (10%).

Table 10: The Pro-Competitive Effects of NAFTA, Domestic Products

	Dependent Variable: $\log \mu_{ijt}$	
	(1)	(2)
τ_{it}^{output}	0.11^a	0.12^a
	(.04)	(.04)
τ_{it}^{US}	-0.06	0.00
	(.14)	(.13)
Instruments	No	Yes
First Stage F	-	44.08
Within R^2	0.82	0.83
N	145,887	124,971

Notes: The dependent variable in Columns 1–2 is the log of the markup. Both specifications include third-order polynomials on log marginal costs (coefficients not reported). In Column 2, we instrument the marginal cost polynomial using its lag and intermediate input tariffs. The regressions exclude outliers in the top and bottom 1% of the markup distribution within each sector. Regressions include plant–product fixed effects and sector–year fixed effects using data for the sample of domestic products and all years (1994–2008). Standard errors are clustered at the product level.
Significance: a (1%), b (5%), and c (10%).

Table 11: The Pro-Competitive Effects of NAFTA, Exported Products

Dependent Variable: $\log \mu_{ijt}$		
	(1)	(2)
τ_{it}^{output}	0.02	-0.01
	(.08)	(.08)
τ_{it}^{US}	-0.54^c	-0.55^c
	(.33)	(.33)
Instruments	No	Yes
First Stage F	-	7.66
Within R^2	0.88	0.88
N	$28,738$	$23,207$

Notes: The dependent variable in Columns 1–2 is the log of the markup. Both specifications include third-order polynomials on log marginal costs (coefficients not reported). In Column 2, we instrument the marginal cost polynomial using its lag and intermediate input tariffs. The regressions exclude outliers in the top and bottom 1% of the markup distribution within each sector. Regressions include plant–product fixed effects and sector–year fixed effects using data for the sample of exported products and all years (1994–2008). Standard errors are clustered at the product level.
Significance: a (1%), b (5%), and c (10%).

Table A1: The Impact of Tariffs on Prices, Marginal Costs, and Markups
Domestic Products by Import Status

	$\log P_{ijt}$	$\log MC_{ijt}$	$\log \mu_{ijt}$
	(1)	(2)	(3)
τ_{it}^{output}	0.13^c	0.11	0.02
	(.07)	(.16)	(.16)
τ_{jt}^{input}	0.11^c	$-.13$	0.24
	(.06)	(.19)	(.18)
τ_{it}^{US}	0.13	1.07^a	-0.93^a
	(.16)	(.37)	(.37)
$Import \times \tau_{it}^{output}$	-0.01	0.08	-0.09
	(.06)	(.16)	(.16)
$Import \times \tau_{jt}^{input}$	0.02	0.35^a	-0.33^b
	(.04)	(.13)	(.14)
$Import \times \tau_{it}^{US}$	-0.17	0.07	-0.24
	(.14)	(.38)	(.38)
Within R^2	0.07	0.01	0.02
N	145,887	145,887	145,887

Notes: Dependent variables in Columns 1–3 are the logs of prices, marginal costs, and markups, respectively. The regressions exclude outliers in the top and bottom 1% of the markup distribution within each sector. Regressions include plant–product and sector–year fixed effects using data for the sample of domestic products and all years (1994–2008). Standard errors are clustered at the product level.
Significance: a (1%), b (5%), and c (10%).

Table A2: The Impact of Tariffs on Prices, Marginal Costs, and Markups
Domestic Products by Export Status

	$\log P_{ijt}$	$\log MC_{ijt}$	$\log \mu_{ijt}$
	(1)	(2)	(3)
τ_{it}^{output}	0.16^a	0.08^c	0.08
	(.05)	(.04)	(.12)
τ_{jt}^{input}	0.09^c	-0.06	0.16
	(.05)	(.18)	(.18)
τ_{it}^{US}	0.03	0.94^a	-0.91^a
	(.14)	(.34)	(.35)
$Export \times \tau_{it}^{output}$	-0.07	0.22	-0.29^c
	(.05)	(.16)	(.16)
$Export \times \tau_{jt}^{input}$	0.11^b	0.52^a	-0.42^a
	(.05)	(.17)	(.17)
$Export \times \tau_{it}^{US}$	0.13	0.40	-0.27
	(.15)	(.42)	(.42)
Within R^2	0.07	0.01	0.02
N	145,887	145,887	145,887

Notes: Dependent variables in Columns 1–3 are the logs of prices, marginal costs, and markups, respectively. The regressions exclude outliers in the top and bottom 1% of the markup distribution within each sector. Regressions include plant–product and sector–year fixed effects using data for the sample of domestic products and all years (1994–2008). Standard errors are clustered at the product level.
Significance: a (1%), b (5%), and c (10%).

Table A3: The Competitive Effects of NAFTA
Domestic Products by Classification

	Homogeneous		Non-Homogeneous	
	(1)	(2)	(3)	(4)
τ_{it}^{output}	0.78^{c}	0.96^{c}	0.07^{c}	0.07^{c}
	(.40)	(.56)	(.04)	(.04)
τ_{it}^{US}	0.49	1.24	-0.03	-0.00
	(1.19)	(1.28)	(.14)	(.14)
Instruments	No	Yes	No	Yes
First Stage F	-	12.33	-	40.63
Within R^2	0.86	0.86	0.81	0.83
N	$8,721$	$7,591$	$136,228$	$116,600$

Notes: The dependent variable in Columns 1–4 is the log of the markup. All specifications include third-order polynomials on log marginal costs (coefficients not reported). In Columns 2 and 4, we instrument the marginal cost polynomial using its lag and intermediate input tariffs. The regressions exclude outliers in the top and bottom 1% of the markup distribution within each sector. Regressions include plant–product fixed effects and sector–year fixed effects using data for the sample of domestic products and all years (1994–2008). Standard errors are clustered at the product level.
Significance: a (1%), b (5%), and c (10%).

Table A4: The Competitive Effects of NAFTA
Exported Products by Classification

	Homogeneous		Non-Homogeneous	
	(1)	(2)	(3)	(4)
τ_{it}^{output}	0.83	0.61	-0.02	-0.03
	(.80)	(.77)	(.07)	(.08)
τ_{it}^{US}	1.05	0.56	-0.67^{c}	-0.64^{c}
	(.70)	(.78)	(.36)	(.35)
Instruments	No	Yes	No	Yes
First Stage F	-	8.27	-	47.36
Within R^2	0.88	0.86	0.88	0.88
N	$1,333$	$1,045$	$27,405$	$22,162$

Notes: The dependent variable in Columns 1–4 is the log of the markup. All specifications include third-order polynomials on log marginal costs (coefficients not reported). In Columns 2 and 4, we instrument the marginal cost polynomial using its lag and intermediate input tariffs. The regressions exclude outliers in the top and bottom 1% of the markup distribution within each sector. Regressions include plant–product fixed effects and sector–year fixed effects using data for the sample of exported products and all years (1994–2008). Standard errors are clustered at the product level.
Significance: a (1%), b (5%), and c (10%).

Table A5: The Competitive Effects of NAFTA
Domestic Products by Size

	Large		Small	
	(1)	(2)	(3)	(4)
τ_{it}^{output}	0.12^a	0.13^a	0.12	0.15
	(.04)	(.04)	(.09)	(.10)
τ_{it}^{US}	-0.14	-0.10	0.26	0.27
	(.16)	(.16)	(.25)	(.22)
Instruments	No	Yes	No	Yes
First Stage F	-	50.82	-	22.35
Within R^2	0.85	0.85	0.84	0.85
N	$73,612$	$64,455$	$75,261$	$61,929$

Notes: The dependent variable in Columns 1–4 is the log of the markup. All specifications include third-order polynomials on log marginal costs (coefficients not reported). In Columns 2 and 4, we instrument the marginal cost polynomial using its lag and intermediate input tariffs. The regressions exclude outliers in the top and bottom 1% of the markup distribution within each sector. Regressions include plant–product fixed effects and sector–year fixed effects using data for the sample of domestic products and all years (1994–2008). Large plants are those with more than 150 employees. Standard errors are clustered at the product level. Significance: a (1%), b (5%), and c (10%).

Table A6: The Competitive Effects of NAFTA
Exported Products by Size

	Large		Small	
	(1)	(2)	(3)	(4)
τ_{it}^{output}	-0.05	-0.07	0.15	0.17
	(.05)	(.06)	(.15)	(.14)
τ_{it}^{US}	-0.39	-0.40	-2.34^a	-2.19^a
	(.32)	(.41)	(.09)	(.08)
Instruments	No	Yes	No	Yes
First Stage F	-	10.20	-	8.98
Within R^2	0.89	0.85	0.86	0.86
N	$21,734$	$17,902$	$7,600$	$5,572$

Notes: The dependent variable in Columns 1–4 is the log of the markup. All specifications include third-order polynomials on log marginal costs (coefficients not reported). In Columns 2 and 4, we instrument the marginal cost polynomial using its lag and intermediate input tariffs. The regressions exclude outliers in the top and bottom 1% of the markup distribution within each sector. Regressions include plant–product fixed effects and sector–year fixed effects using data for the sample of exported products and all years (1994–2008). Large plants are those with more than 150 employees. Standard errors are clustered at the product level. Significance: a (1%), b (5%), and c (10%).

Table A7: The Competitive Effects of NAFTA
Domestic Products by Initial Markup

	(1)	(2)
τ_{it}^{output}	0.08^b	0.11^b
	(.04)	(.04)
τ_{it}^{US}	-0.06	-0.01
	(.14)	(.04)
$High\_Tariff_{0,ijt} \times \tau_{it}^{output}$	0.10^c	0.07
	(.06)	(.05)
$High\_Tariff_{0,ijt} \times \tau_{it}^{US}$	0.24	0.32
	(.21)	(.23)
Instruments	No	Yes
First Stage F	-	39.70
Within R^2	0.84	0.85
N	$148,873$	$127,341$

Notes: The dependent variable in Columns 1 and 2 is the log of the markup. Both specifications include third-order polynomials on log marginal costs (coefficients not reported). In Column 2, we instrument the marginal cost polynomial using its lag and intermediate input tariffs. The regressions exclude outliers in the top and bottom 1% of the markup distribution within each sector. Regressions include plant–product fixed effects and sector–year fixed effects using data for the sample of domestic products and all years (1994–2008). Standard errors are clustered at the product level.
Significance: a (1%), b (5%), and c (10%).

Table A8: The Competitive Effects of NAFTA
Exported Products by Initial Markup

	(1)	(2)
τ_{it}^{output}	-0.00	-0.03
	(.06)	(.06)
τ_{it}^{US}	-0.68^{c}	-0.60
	(.40)	(.44)
$High\_Tariff_{0,ijt} \times \tau_{it}^{output}$	0.86	0.23
	(.53)	(.55)
$High\_Tariff_{0,ijt} \times \tau_{it}^{US}$	-0.05	0.04
	(.86)	(1.04)
Instruments	No	Yes
First Stage F	-	9.42
Within R^2	0.89	0.88
N	$28,738$	$23,207$

Notes: The dependent variable in Columns 1 and 2 is the log of the markup. Both specifications include third-order polynomials on log marginal costs (coefficients not reported). In Column 2, we instrument the marginal cost polynomial using its lag and intermediate input tariffs. The regressions exclude outliers in the top and bottom 1% of the markup distribution within each sector. Regressions include plant–product fixed effects and sector–year fixed effects using data for the sample of exported products and all years (1994–2008). Standard errors are clustered at the product level.
Significance: a (1%), b (5%), and c (10%).

Table A9: The Impact of NAFTA on Product Introduction

	Contemporaneous			Lagged		
	(1)	(2)	(3)	(4)	(5)	(6)
τ_{jt}^{output}	0.000	0.000	−0.001*	0.001	0.001	−0.001
	(.001)	(.001)	(.001)	(.001)	(.001)	(.001)
τ_{jt}^{input}	0.029***	0.029***	0.023**	0.028***	0.027***	0.021***
	(.003)	(.003)	(.003)	(.003)	(.003)	(.003)
τ_{jt}^{US}	0.007	0.007	0.009*	0.004	0.003	0.006
	(.005)	(.005)	(.005)	(.004)	(.004)	(.004)
$\tau_{jt}^{output} \times EXP_{jt}$	−0.003*	−0.003	−0.004***	−0.005***	−0.005***	−0.006***
	(.002)	(.002)	(.002)	(.002)	(.002)	(.002)
$\tau_{jt}^{input} \times EXP_{jt}$	−0.008**	−0.008**	−0.006	−0.008*	−0.008*	−0.005
	(.004)	(.004)	(.004)	(.004)	(.004)	(.004)
$\tau_{jt}^{US} \times EXP_{jt}$	−0.024***	−0.021***	−0.018**	−0.015**	−0.012*	−0.010*
	(.008)	(.008)	(.008)	(.007)	(.007)	(.006)
Plant Controls	✓	✓	✓	✓	✓	✓
Year FE	✓	✓		✓	✓	
State FE		✓	✓		✓	✓
Sector–Year FE			✓			✓
Pseudo R^2	0.069	0.071	0.077	0.027	0.030	0.036
N	54,091	54,091	54,091	48,571	48,571	48,571

Notes: Columns 1–6 gives the results of the probit specification. Columns 1–3 include the specification with the independent variables in their contemporaneous values, and Columns 4–6 show the results with the independent variables lagged one year. All specifications include labor, capital, and material expenditures, as well as import and export status as plant-level controls. The coefficients on controls are omitted for brevity. Regressions include plant-level for the years 1994–2008.
Significance: a (1%), b (5%), and c (10%).

Table A10: The Impact of NAFTA on TFPQ

	(1)	(2)	(3)	(4)	(5)
τ_{jt}^{output}	-1.272^{***}	-1.225^{***}	-1.255^{***}	-1.363^{***}	-0.200^{***}
	(.224)	(.218)	(.221)	(.245)	(.056)
τ_{jt}^{input}	-2.307^{***}	-2.453^{***}	-2.197^{***}	-1.868^{*}	-2.024^{***}
	(.467)	(.466)	(.781)	(1.001)	(.166)
τ_{jt}^{US}	-4.009^{***}	-3.802^{***}	-3.383^{***}	-1.824^{*}	-0.550^{**}
	(.769)	(.762)	(.852)	(1.073)	(.264)
Plant Controls	✓	✓	✓	✓	✓
State FE		✓	✓	✓	
Year FE			✓		
Sector–Year FE				✓	✓
Plant FE					✓
R^2	0.072	0.083	0.084	0.107	0.088
N	37,498	37,498	37,498	37,498	37,498

Notes: The dependent variable in Columns 1–5 is the logarithm of a plant's productivity. The specifications in Columns 1–5 vary only on the set of fixed effects included in the estimation. The R^2 reported in Column 5 corresponds to the within R^2. All specifications include import and export status, total industry sales excluding the plant's sales, and the HHI index of concentration as plant-level controls. The coefficients on controls are omitted for brevity. Regressions include plant-level for the years 1994–2008. Standard errors are clustered at the plant level.
Significance: a (1%), b (5%), and c (10%).

Figure 1: Mexican Tariff Schedule Under NAFTA, Selected Products

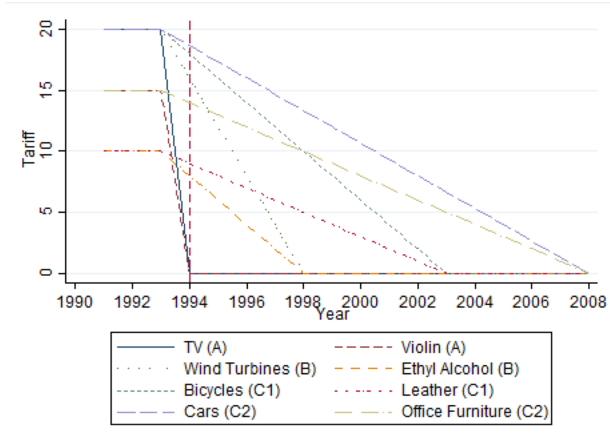

Notes: The figure plots Mexico's tariff decline schedule for a selected sample of products in different groups. Products in group A had to have their tariff phased out by the year 1994, those in group B by 1998, those in group C1 by 2003 and finally products in group C2 had to be phased out in 2008.
Source: NAFTA Agreement, Annex 302.2 Section B.

Figure 2: Average Tariffs (1993–2008)

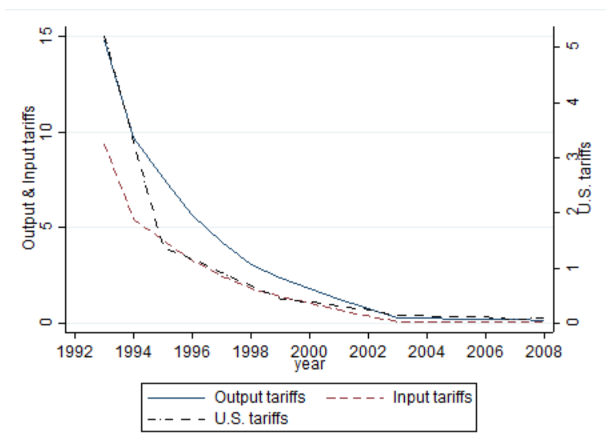

Notes: The figure plots the averages of our measures of output tariffs, intermediate input tariffs, and U.S. tariffs for the years 1993–2008. Output and intermediate input tariffs on the left scale, and U.S. tariffs on the right scale.

Figure 3: Quality and Average Wages

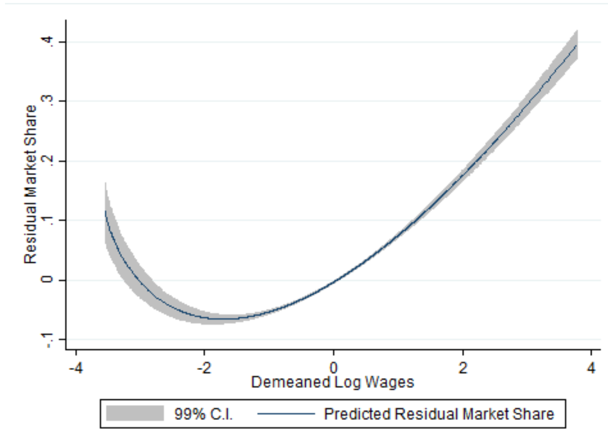

Notes: The figure plots the the best-fitted polynomial of residuals from a regression of product market shares on prices and product dummies (y-axis) and the log of average wages demeaned by product–market fixed effects (x-axis) for the full 1994–2008 sample. Average wages were constructed by dividing total wage bill by total number of employees. The shaded area indicates a 99% confidence interval.

Figure 4: Average and Implied Material Expenditure Shares

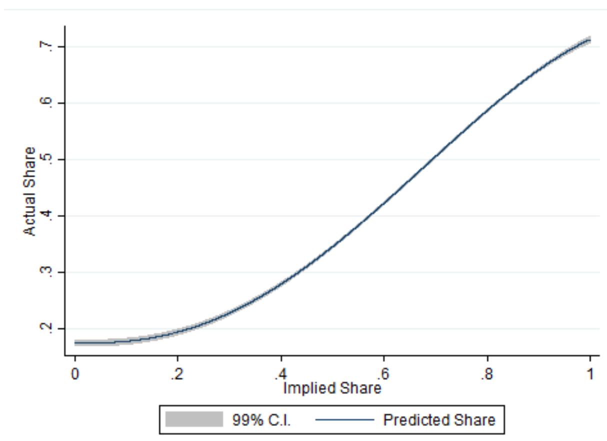

Notes: The figure shows the best-fitted polynomial of the share of materials expenditure out of total expenditures in labor, capital, and materials (y-axis), and the materials expenditure share implied by the estimated elasticities (x-axis). The shaded area indicates a 99% confidence interval.

Figure 5: Estimated Markups and Accounting Revenue/Variable Costs

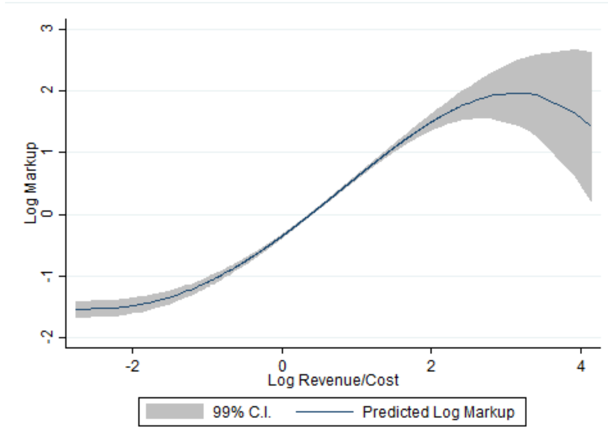

Notes: The figure shows the best-fitted polynomial of the logarithm of estimated markups (y-axis) and the log of the ratio of accounting revenue over variable costs (x-axis). The shaded area indicates a 99% confidence interval. The figure excludes outliers below the 1st and above the 99th percentiles of the markup distribution in each sector.

Figure 6: Markups, Marginal Costs and Quantity Produced

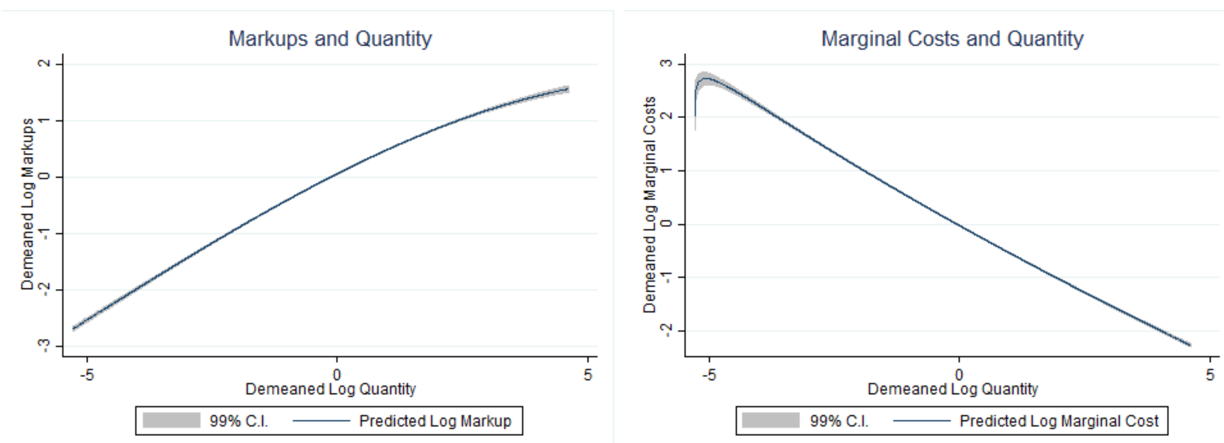

Notes: The figure shows the best-fitted polynomials of the logarithm of markups and marginal costs, demeaned by product–market fixed effects (y-axis) and the demeaned logarithm of quantity produced (x-axis). The shaded area indicates a 99% confidence interval. The figure excludes outliers below the 1st and above the 99th percentiles of the markup distribution in each sector.

Figure 7: The Impact of Output Tariffs on P, MC, and μ

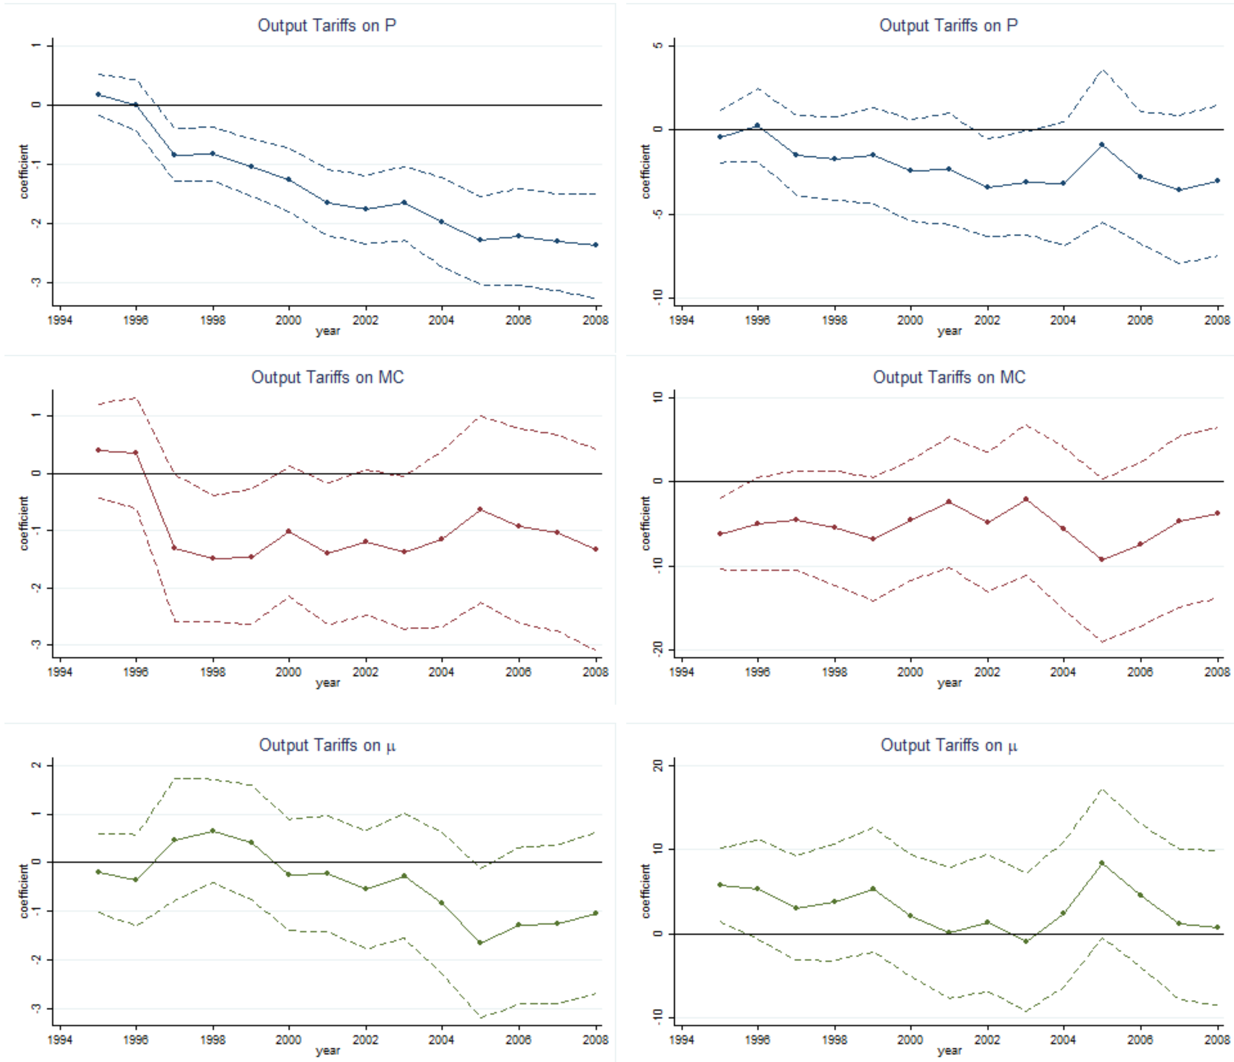

(a) Domestic Products (b) Exported Products

Notes: Each point in the graph represents an individual regression coefficient multiplied by -1. The dashed lines show the 95% confidence intervals.

Figure 8: The Impact of Input Tariffs on P, MC, and μ

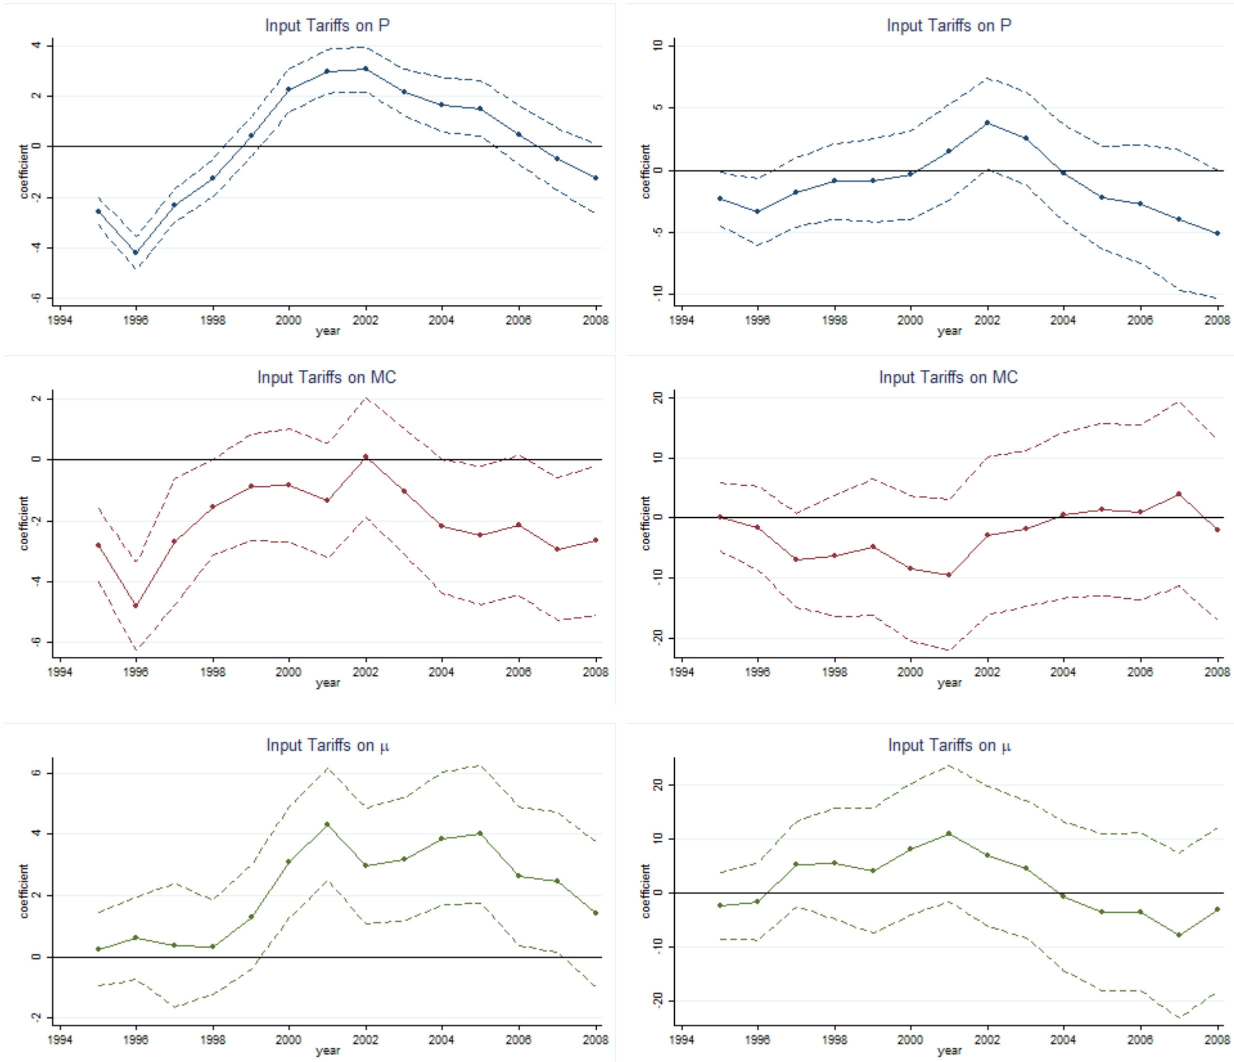

(a) Domestic Products (b) Exported Products

Notes: Each point in the graph represents an individual regression coefficient multiplied by -1. The dashed lines show the 95% confidence intervals.

Figure 9: The Impact of Input Tariffs on Quality

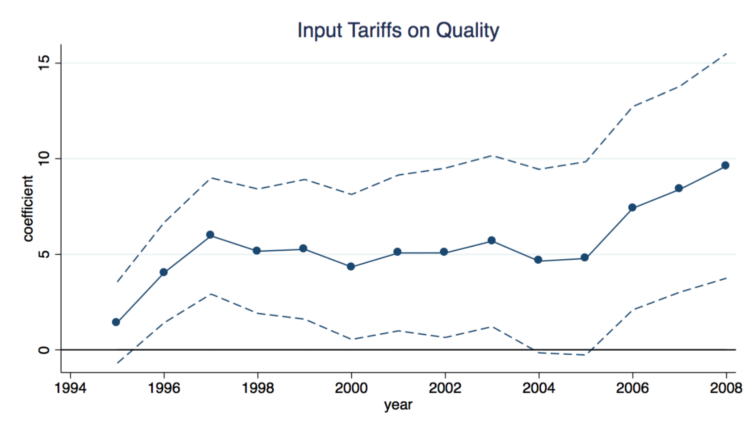

(a) Domestic Products

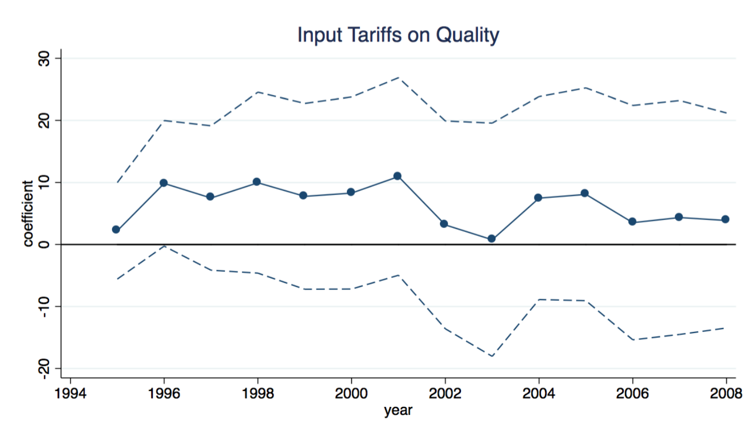

(b) Exported Products

Notes: Each point in the graph represents an individual regression coefficient multiplied by -1. The dashed lines show the 95% confidence intervals.

Figure 10: The Impact of U.S. Tariffs on P, MC, and μ

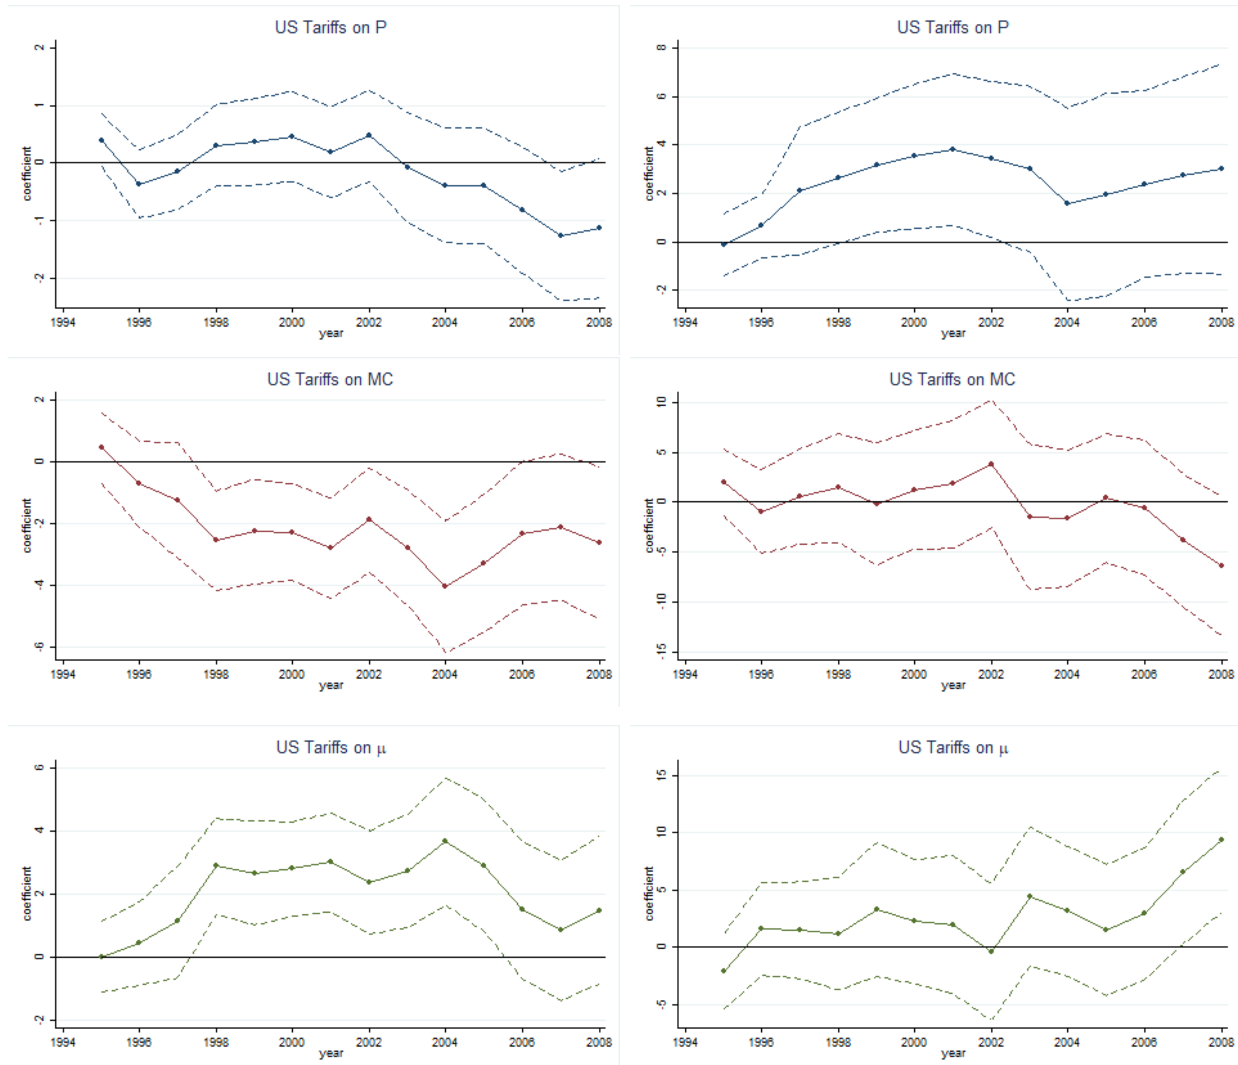

(a) Domestic Products

(b) Exported Products

Notes: Each point in the graph represents an individual regression coefficient multiplied by -1. The dashed lines show the 95% confidence intervals.