

CAMBRIDGE WORKING PAPERS IN ECONOMICS
JANEWAY INSTITUTE WORKING PAPERS

The Pro-competitive Effects of Trade Agreements

Meredith A.
Crowley
University of
Cambridge

Lu
Han
University of
Liverpool

Thomas
Prayer
University of
Cambridge

Abstract

How does trade policy affect competition? Using the universe of product exports by firms from eleven low and middle-income countries, we document that tariff reductions under trade agreements have strong procompetitive effects – they encourage entry and reduce the (tariff exclusive) price-cost markups of exporters. This finding, that markups fall with tariff cuts, contradicts a core prediction of standard oligopolistic competition models of trade. We extend a workhorse international pricing model of oligopolistic competition to include multiple countries and a rich preference structure. Our preference structure allows for fierce competition among firms from the same country and less intense competition among firms from different countries. We show a firm’s optimal markup after a tariff cut can rise or fall depending on the parameters of the preference structure and tariff-induced reallocation of market share among firms and across countries.

Reference Details

2240 Cambridge Working Papers in Economics
2216 Janeway Institute Working Paper Series

Published 30 June 2022
Revised 11 July 2022

Key Words trade agreements, variable markups, markup elasticity, trade elasticity, competition policy, firm level data

JEL Codes F13, F14, F15

Websites www.econ.cam.ac.uk/cwpe
www.janeway.econ.cam.ac.uk/working-papers

The Pro-competitive Effects of Trade Agreements

Meredith A. Crowley^{*} Lu Han[†] Thomas Prayer[‡]

11 July 2022

Abstract

How does trade policy affect competition? Using the universe of product exports by firms from eleven low and middle-income countries, we document that tariff reductions under trade agreements have strong procompetitive effects – they encourage entry and *reduce* the (tariff exclusive) price-cost markups of exporters. This finding, that markups fall with tariff cuts, contradicts a core prediction of standard oligopolistic competition models of trade. We extend a workhorse international pricing model of oligopolistic competition to include multiple countries and a rich preference structure. Our preference structure allows for fierce competition among firms from the same country and less intense competition among firms from different countries. We show a firm’s optimal markup after a tariff cut can rise or fall depending on the parameters of the preference structure and tariff-induced reallocation of market share among firms and across countries.

^{*}University of Cambridge, Janeway Institute Cambridge, CEPR, and the UK in a Changing Europe; email: meredith.crowley@econ.cam.ac.uk.

[†]University of Liverpool, Janeway Institute Cambridge and CEPR; email: hanlulong@gmail.com

[‡]University of Cambridge; email: tp392@cam.ac.uk.

JEL classification: F13, F14, F15

Keywords: trade agreements, variable markups, markup elasticity, trade elasticity, competition policy, firm level data.

Acknowledgements: We are grateful to Mary Amiti, Costas Arkolakis, Andy Bernard, Tibor Besedes, Ariel Burstein, Hector Calvo Pardo, Giancarlo Corsetti, Vasco Carvalho, Matt Elliott, Eric French, Michael Irlacher, Glenn Magerman, Emanuel Ornelas, Mathieu Parenti, Bob Staiger, Katheryn Russ, Steve Redding, David Weinstein and Yoto Yotov for constructive feedback and thoughtful questions that helped us to improve this paper. We would also like to thank participants of the NBER Conference on Trade and Trade Policy in the 21st Century (Boston), the Royal Economic Society 2022 Annual Conference, the Trade Policy and Global Value Chains Workshop at the University of Mannheim, the UK in a Changing Europe Conference on The Economics of Brexit (London), the IAAE 2022 Annual Conference (London), the Midwest Trade and Theory 2022 Conference (Purdue), and the Universite Libre de Bruxelles/ECARES trade seminar for engaging discussions. Ana Margarida Fernandes and Fujie Wang provided access to the Exporter Dynamics Database used in this study and answered our many questions about these data, for which we are grateful. This project developed out of a paper created for the World Bank's Deep Trade Agreements Seminar, and we would like to thank Jeff Bergstrand and Paulo Bastos, our discussants, for raising thought-provoking questions that prompted us to write this paper. We are indebted to the Economic and Social Research Council of the United Kingdom/UK in a Changing Europe Senior Fellowship ES/T000732/1 for financial support that made this research possible.

1 Introduction

Do tariff liberalizations under trade agreements impact the structure of markets and competition among firms? How does trade policy affect exporters’ price-cost markups? The answers to these questions have important implications for welfare. A recent contribution from [Arkolakis, Costinot, Donaldson and Rodríguez-Clare \(2018\)](#) shows that, in a large class of trade models, the pro-competitive gains of trade liberalizations depend crucially on exporters’ markup adjustments. If exporters raise their markups in response to a preferential tariff cut, pro-competitive gains from trade are “elusive” because potential welfare gains due to markup reductions by domestic firms will be offset by equivalent or larger welfare losses due to markup increases by foreign exporters. While several papers have confirmed that domestic firms do in fact reduce their markups in response to a trade liberalization, empirical evidence on how foreign exporters adjust their markups remains scarce.¹

In this paper, we investigate exporters’ markup responses to trade policy changes for eleven low and middle income economies by integrating their annual customs records with information on 83 preferential trade agreements and data on bilateral import tariffs for 165 destinations.² The unique structure of this multi-origin panel allows us to identify changes in markups and market shares by exploiting variation in firms’ product-level export unit values (i.e., prices) and export sales across destinations and over time. What we find is puzzling. While a 10% reduction in bilateral tariffs raises a foreign exporter’s market share in a destination by 7.8%, it *lowers* an exporter’s pre-tariff markup by 4.1%. The first finding is consistent with classical Vinerian trade diversion: preferred origins gain market share as a result of trade liberalizations. It is the second finding which presents the puzzle. In the classic literature on trade under imperfect competition ([Brander and Krugman \(1983\)](#), [Helpman and Krugman \(1985\)](#), [Eaton and Grossman \(1986\)](#) and [Markusen and Venables \(1988\)](#)) and the more recent international macro literature on trade under oligopolistic competition with endogenously variable markups ([Atkeson and Burstein \(2008\)](#), [Edmond, Midrigan and Xu \(2015\)](#)), an exporting firm’s markup moves in the same direction as its market share.³ Thus,

¹For example, recent work by [De Loecker, Goldberg, Khandelwal and Pavcnik \(2016\)](#) and [Edmond, Midrigan and Xu \(2015\)](#) has found that trade liberalizations reduce the prices charged by domestic firms. Several papers ([Bown and Crowley \(2006\)](#), [Amiti, Redding and Weinstein \(2019\)](#), and [Fajgelbaum, Goldberg, Kennedy and Khandelwal \(2020\)](#)) examine foreign unit value responses to trade policy changes, but their product-level datasets do not allow for an analysis of markups. A recent study by [Kikkawa, Mei and Santamarina \(2019\)](#) uses survey data on Mexican firms to examine the impact of NAFTA on markups domestically and for exported products.

²The eleven countries are: Albania (2004-2012); Bulgaria (2001-2006); Burkina Faso (2005-2012); China (2000-2006); Egypt (2005-2013); Malawi (2006-2012); Mexico (2000-2012); Peru (2000-2013); Senegal (2000-2012); Uruguay (2001-2012); and Yemen (2008-2012).

³Exact predictions are model-specific and sensitive to assumptions about market structure. For example, [Markusen and Venables \(1988\)](#) map out assumptions about market segmentation and free entry to show

a tariff cut that *raises* a firm’s market share in a destination should *raise* the firm’s (tariff exclusive) markup.

Our theoretical contribution is to introduce a more general framework for analysing international trade that can explain this empirical puzzle. We extend [Atkeson and Burstein \(2008\)](#)’s seminal contribution on international pricing in two *empirically relevant* directions. First, we extend their original two-country model into a multi-country framework allowing for oligopolistic competition among firms from multiple origins in multiple destinations. Second, we introduce a third nest into [Atkeson and Burstein \(2008\)](#)’s two-tiered CES preference structure, thereby allowing for a higher degree of substitutability among goods produced by firms from the same origin than among goods produced by firms from different origins. This second generalization implies that competition among firms from the same origin country can be fiercer than that among firms from different origins. Under this more general framework, a firm’s markup adjustment to a tariff cut depends on its elasticity of demand, which, in turn, depends on the market structure in the destination and the firm’s relative market power. The precise elasticity of demand facing a firm is determined by two different market share measures and three elasticities of substitution; it nests the demand elasticities in [Atkeson and Burstein \(2008\)](#) and [Melitz \(2003\)](#) as special cases.

Using this framework, we decompose the conventional relationship between changes in market shares and price-cost markups in response to a preferential tariff cut into contributions from two separate market share reallocation effects. The first is an “across-origin” reallocation effect that captures how the increase in market access for the origin country as a whole affects origin firms’ market power, and hence, markups. The second is a “within-origin” reallocation effect that captures how a change in an individual firm’s share of its origin’s exports affects the firm’s markup. We show that the two reallocation effects place opposing pressures on markups. On the one hand, the “across-origin” reallocation effect increases the optimal markup of exporters from the preferred origin. On the other hand, lower tariffs encourage entry from a preferred origin, which reduces the “within-origin” market share of incumbent exporters from that preferred origin, lowering their optimal markups. Depending on the magnitude of these two reallocation effects, exporters may either raise (anti-competitive) or lower (pro-competitive) their markups. Theoretically, we show changes in the “within-origin” reallocation channel tend to dominate if goods produced by firms from

changes in assumptions about free entry nullify or flip predictions arising from changes in trade policy. That said, it is relatively uncontroversial to say that in models where firms face a downward-sloping import demand curve, we would expect the benefits of a tariff reduction to be split between consumers and exporters. Part of the tariff cut should be passed on to consumers in the form of lower consumer prices and part should be passed onto exporters as an *increase* in the tariff-exclusive export price, as a result of an increase in the price-cost markups of individual firms.

the same origin are more substitutable with each other than with those produced by firms from different origins.

Our theoretical explanation of the markup puzzle relies on two mechanisms that we can verify empirically. First, our model predicts that the origin’s market share in a destination *rises* while each individual firm’s within-origin market share *falls* in response to a preferential tariff cut. And indeed, we estimate that a 10% preferential tariff liberalization leads to a 23% increase in an origin country’s market share in the liberalizing destination and a 29% *reduction* in a firm’s share of its origin’s trade with the destination. Second, our model predicts a strong entry effect after a bilateral tariff reduction. Empirically, we find a 10% reduction in bilateral tariffs induces entry from the preferred origin and increases exporter participation by 22%. Taken altogether, the empirical evidence indicates the within-origin market share reallocation effect dominates the across-origin reallocation effect, leading to an average 4.1% decrease in exporting firms’ (pre-tariff) markups.

Our conceptual framework also offers important insights into the relationship between the elasticity of markups to trade policy and the degree of product differentiation. We show that the difference between the “within-origin” and the “across-origin” reallocation effects is more pronounced, and markup adjustments in response to a bilateral tariff reduction are thus larger, for goods whose degree of substitutability in consumer demand among firms within an origin shows more substantial differences relative to that among firms from different origins. Our empirical findings provide indirect evidence for this prediction. Markup adjustments for more differentiated products and consumption goods are significantly larger than those for homogenous and intermediate goods. Intuitively, homogeneous and intermediate goods sold by firms from different origins are quite similar. For these goods, the within-origin and across-origin reallocation effects have similar magnitudes but opposite signs, resulting in little or no overall markup adjustment. In contrast, differentiated and consumption goods offer more scope for differences in the degree of competition within and across origins, resulting in much bigger markup adjustments when tariffs change.

Literature review. Our empirical research builds on a methodologically diverse body of work examining how prices and markups change in response to trade policy changes (Konings and Vandebussche (2005), Bown and Crowley (2006), Amiti and Konings (2007), Pierce (2011), Edmond, Midrigan and Xu (2015), De Loecker, Goldberg, Khandelwal and Pavcnik (2016), Fitzgerald and Haller (2018), Amiti, Redding and Weinstein (2019), Kikkawa, Mei and Santamarina (2019), Fajgelbaum, Goldberg, Kennedy and Khandelwal (2020)) and exchange rate movements (Fitzgerald and Haller (2014), De Blas and Russ (2015), Amiti, Itshhoki and Konings (2019), Corsetti, Crowley, Han and Song (2021), Corsetti, Crowley and Han (2022)).

A unique feature of our paper compared to previous studies is that we use product-level exports by firms from multiple origins. This has two advantages: (a) it allows us to simultaneously control for two confounding factors that are important for analyzing firms’ markup adjustments, time-varying marginal cost at the firm-product level and time-varying shifts in demand for a product in a destination,⁴ and (b) it allows us to directly map changes in firms’ export sales into the market share measures that determine the elasticity of demand in our theoretical framework. Altogether, it allows us to provide, to the best of our knowledge, the first empirical evidence on how exporters from multiple origins adjust their markups and market shares in a destination market in response to preferential trade policy changes.⁵

The starting point for our theoretical analysis of firms’ responses to trade policy changes is the contribution of [Atkeson and Burstein \(2008\)](#) to the open economy macro literature on international pricing. We extend the [Atkeson and Burstein \(2008\)](#) model and show that exporters may increase or decrease their markups in response to a bilateral tariff cut. This finding relies on two observations. First, within narrowly defined product markets, the number of firms exporting from an origin to a destination is small – the median number in our sample of eleven origin countries is two.⁶ In such cases, entry and exit of exporters can have non-negligible impacts on market structure and incumbent exporting firms’ markups. Second, goods produced by firms from the same origin may be more substitutable with each other than with goods produced by firms from different origins. This implies the within-origin market share changes could have a bigger impact on a firm’s demand elasticity and markup than the across-origin market share changes. Thus, markups can fall even when tariff cuts lead to no change or a modest increase in an exporter’s market share in a destination. A key implication of our model is that, if exporters reduce their markups in response to a bilateral tariff liberalization, the pro-competitive gains from trade will be strictly positive. This contrasts with the predictions of [Arkolakis, Costinot, Donaldson and Rodríguez-Clare \(2018\)](#), who show that pro-competitive welfare gains from small trade liberalizations are “elusive” in a large class of models with variable markups under reasonable calibrations of markup and demand elasticities, and with the analysis of [Edmond, Midrigan and Xu \(2015\)](#), which shows that even in the [Atkeson and Burstein \(2008\)](#) framework, there are no pro-competitive gains from small trade liberalizations.⁷

⁴This is achieved through the use of firm-product-origin-time and product-destination-time fixed effects.

⁵A previous contribution from [Bas, Mayer and Thoenig \(2017\)](#) shows how two firm-level datasets can be employed to provide better estimates of aggregate bilateral trade elasticities, but does not explore prices or markups.

⁶This is the median number of firms conditional on an origin country exporting to a destination. The unconditional median is zero. See table [A8](#) for more details.

⁷This result is interesting as it is obtained for a model which falls outside the set of models considered by [Arkolakis, Costinot, Donaldson and Rodríguez-Clare \(2018\)](#) because it allows for strategic interactions

The rest of the paper is organised as follows. Section 2 introduces our new model. Section 3 presents our data and section 4 our empirical strategy. We discuss our empirical results in section 5 and conclude in section 6.

2 Conceptual framework

In this section, we present a multi-country framework that allows us to study how firms compete and adjust their markups in response to trade liberalizations. We follow Atkeson and Burstein (2008) and Edmond, Midrigan and Xu (2015) and consider a nested CES demand structure with a finite number of producers in each industry.

The world consists of $H \geq 3$ countries and trade among countries is indexed by origin $o \in \mathcal{H}$ and destination $d \in \mathcal{H}$.⁸ In each country, there is a continuum of unit mass of industries, indexed by i , selling tradable goods.⁹ Final consumption Y_{dt} and the price of the final consumption good P_{dt} in each country d in period t are aggregated over industries i :

$$Y_{dt} = \left(\int_i y_{idt}^{\frac{\eta-1}{\eta}} di \right)^{\frac{\eta}{\eta-1}}, \quad P_{dt} = \left(\int_i p_{idt}^{1-\eta} di \right)^{\frac{1}{1-\eta}} \quad (1)$$

where $\eta > 1$ is the elasticity of substitution across industries. Industry-level output y_{idt} and the industry-level price index p_{idt} are obtained by aggregating products across different origins:

$$y_{idt} = \left(\sum_{o \in \mathcal{H}} y_{io dt}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}}, \quad p_{idt} = \left(\sum_{o \in \mathcal{H}} p_{io dt}^{1-\rho} \right)^{\frac{1}{1-\rho}} \quad (2)$$

where $\rho \geq \eta$ is the elasticity of substitution across products from different origins. Within each industry-origin-destination triplet, there is a finite number of firms, each producing a differentiated variety. The industry-origin-destination level output $y_{io dt}$ and price $p_{io dt}$ are obtained by aggregating across firms from the same origin:

$$y_{io dt} = \left(\sum_{f \in \mathcal{F}_{io dt}} (\alpha_{fio dt})^{1/\sigma} y_{fio dt}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \quad p_{io dt} = \left(\sum_{f \in \mathcal{F}_{io dt}} (\alpha_{fio dt}) p_{fio dt}^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \quad (3)$$

where $\sigma \geq \rho$ is the elasticity of substitution across varieties from the same origin, $\alpha_{fio dt}$ is a demand/preference shifter and $\mathcal{F}_{io dt}$ represents the set of active firms that sell product i

and uses a nested CES preference structure.

⁸Throughout our paper, we use calligraphy math symbols to indicate a set of elements.

⁹In our empirical analysis, an “industry” is an HS6 product. We use the words “industry” and “product” interchangeably throughout the paper.

from origin o to destination d at time t .¹⁰

The key difference of our setting compared to a standard two-country [Atkeson and Burstein \(2008\)](#) model is that we introduce an additional layer of aggregation across firms from the same origin (i.e., equation (3)), which allows for a different elasticity of substitution within and across origins, i.e., $\sigma \neq \rho$. For example, a consumer may view the t-shirts produced by two Chinese exporters, which very likely share access to the same specialised resources, operate under the same institutions, and compete with each other in many of the same markets, as very similar to each other and at the same time consider them quite different from a t-shirt made by, say, a Japanese or American firm which may use different fabrics and designs.

Production. Labor is inelastically supplied and immobile across countries, and wages are identical across sectors in a given country. The production function is linear in labour L and productivity Ω , i.e., $Y \equiv \Omega L$. The marginal cost of the firm is thus $mc_{fio t} = W_{ot}/\Omega_{fio t}$, where W_{ot} is the nominal wage of the origin country o at time t and $\Omega_{fio t}$ is the productivity of firm f in industry i from country o at time t .

Price and export decisions. Firms compete by simultaneously choosing whether to enter a market, indicated by $\phi_{fio dt} \in \{0, 1\}$, and their optimal price $p_{fio dt}$ if they enter. Since the production function features constant returns to scale, firms make their pricing and entry decisions separately for each destination market. The profit maximization problem of firm f in industry i from origin o selling in destination d is given by:

$$\pi_{fio dt} = \max_{p_{fio dt}, \phi_{fio dt}} \left[y_{fio dt} \left(\frac{p_{fio dt}}{\tau_{io dt}} - mc_{fio t} \right) - W_{ot} \zeta_x \right] \phi_{fio dt}$$

subject to

$$y_{fio dt} = \alpha_{fio dt} \left(\frac{p_{fio dt}}{p_{io dt}} \right)^{-\sigma} \left(\frac{p_{io dt}}{p_{id t}} \right)^{-\rho} \left(\frac{p_{id t}}{P_{dt}} \right)^{-\eta} Y_{dt} \quad (4)$$

where $\tau_{io dt}$ is the bilateral trade cost including tariffs and ζ_x is a constant per-period export cost in terms of labor units. The firm will enter a market if the potential operating profit $y_{fio dt}(p_{fio dt}/\tau_{io dt} - mc_{fio t})$ is larger than the fixed per-period exporting cost $W_{ot}\zeta_x$.¹¹

Upon entry, the optimal price $p_{fio dt}$ and markup $\mu_{fio dt}$ for an exporter f from origin o to

¹⁰We indicate a variable's level of aggregation in our model by its subscript. The most disaggregated variables have five dimensions, f, i, o, d and t , which stand for firm, industry, origin, destination, and time, respectively.

¹¹The production and price decisions in the domestic market are similarly defined with a smaller fixed cost of operating in the domestic market, $\zeta_h < \zeta_x$.

destination d can be derived as:

$$p_{fiomt} = \mu_{fiomt} m c_{fiomt} \tau_{fiomt}, \quad \mu_{fiomt} = \frac{\varepsilon_{fiomt}}{\varepsilon_{fiomt} - 1} \quad (5)$$

where ε_{fiomt} is the price elasticity of demand. In what follows, we discuss the key implications of our extensions for the firm's optimal markup μ_{fiomt} under different assumptions about competition.

2.1 Market structure, competition, and markups

The way in which firms compete depends on the structure of a market, which is characterized by two sets of statistics: (1) the market share distributions of firms and (2) the substitutability of varieties within an origin, across origins and across industries.

The general functional form of the demand elasticity under the triple-nested demand structure described by expressions in (1) - (3) can be derived as follows:¹²

$$\varepsilon_{fiomt} = \sigma - m s_{fiomt} [\sigma - \rho + (\rho - \eta) m s_{io' dt}] \quad (6)$$

where the first market share $m s_{fiomt}$ captures the importance of the firm among all exporters from the origin and the second market share $m s_{io' dt}$ captures the importance of the origin country in the destination market:

$$\underbrace{m s_{fiomt} = \frac{p_{fiomt} y_{fiomt}}{\sum_{f' \in \mathcal{F}_{fiomt}} p_{f' iomt} y_{f' iomt}}}_{\text{firm's within-origin market share}}, \quad \underbrace{m s_{io' dt} = \frac{p_{io' dt} y_{io' dt}}{\sum_{o' \in \mathcal{H}} p_{io' dt} y_{io' dt}}}_{\text{origin's market share in the destination}} \quad (7)$$

In what follows, we show that equation (6) is a generalization, which nests many important models in the literature.

Monopolistic competition. First, there are two important cases where our model converges to a Melitz (2003) model: (a) when the number of firms from the same origin is large enough, e.g., $|\mathcal{F}_{fiomt}| \rightarrow \infty$, and/or (b) when the degree of substitutability is the same for all products, i.e., $\sigma = \rho = \eta$.

In either case, firms compete under monopolistic competition and charge constant markups:

$$\frac{\varepsilon_{fiomt}}{\varepsilon_{fiomt} - 1} = \frac{\sigma}{\sigma - 1} \quad (8)$$

A key implication of this market structure is that the optimal markup is the same across

¹²See Appendix B.1 for the complete derivation.

big (more productive) and small (less productive) firms. In this case, firms will fully pass through any change in tariffs or other trade costs to the consumer price. Both (a) and (b) are strong theoretical assumptions which generate predictions that are not supported in the data. This has led many researchers to turn their attention to models featuring variable markups (Bernard, Eaton, Jensen and Kortum (2003), Melitz and Ottaviano (2008), Atkeson and Burstein (2008), Edmond, Midrigan and Xu (2015), Amiti, Itskhoki and Konings (2019)) with the class of models introduced by Atkeson and Burstein (2008) proving especially useful for studying pricing under oligopolistic competition at the industry level.

Oligopolistic competition at the industry level. Second, our model converges to that of Atkeson and Burstein (2008) if the number of firms operating in an industry is finite and the substitutability of products from different origins is the same, i.e., $\sigma = \rho$.

Under this market structure, the firm will internalize the impact of its competitors' prices at the industry level and the demand elasticity in (6) can be simplified to

$$\varepsilon_{fiotd} = \rho - (\rho - \eta)\omega_{fiotd} \quad (9)$$

where $\omega_{fiotd} = ms_{ioidt}ms_{fiotd}$ is the firm's market share in the destination, capturing the importance of the firm in industry i to destination d at time t . A crucial implication of (9) is that a tariff reduction increases the market shares of firms from the preferred origin in the destination and thus leads to increases in their markups.

Oligopolistic competition among firms from the same origin. If the number of firms from an origin selling a specific product to a particular destination is finite and small but the number of competitors from other origins is large, the firm may view other firms from the same origin as its key competitors and endogenize its impact on the origin-specific industry price index in the destination p_{ioidt} but not on the overall industry price index in the destination p_{idt} . As $ms_{ioidt} \rightarrow 0$, the demand elasticity converges to:

$$\varepsilon_{fiotd} \rightarrow \sigma - ms_{fiotd}(\sigma - \rho) \quad (10)$$

A key feature of (10) is that firms will only adjust their markups according to the level of competition from their peers from the same origin. Contrary to the prediction of the Atkeson and Burstein (2008) case, a tariff reduction will lead to a drop in the average markup of continuing firms from the origin. Since the tariff reduction makes firms from the origin more competitive, some small firms will find it optimal to export, and hence, enter the market, which reduces the average market share ms_{fiotd} of existing firms. The drop in market share in turn increases the demand elasticity, which leads to a reduction in the average markup. While intuitive, our model is the first to formally characterize this oligopolistic competition

at the level of origin, to the best of our knowledge. Next, we show how to build on this intuition and construct a more general model where firms can compete oligopolistically both within and across origins within an industry.

A more general case: oligopolistic competition within both origin and industry. If, in a more general case, we allow for a small number of competitors from the same origin as well as from other countries, the firm will likely view both sets of firms as its competitors. Thus, when the degree of substitutability for varieties produced within the same origin is different from that for varieties produced in different origins (i.e., $\sigma \neq \rho$), the firm will endogenize its impact on both the origin-industry price index p_{iodt} and the industry price index p_{idt} so that its demand elasticity takes the general form characterized in (6).

In this more general case, a preferential tariff reduction will lead to two competing channels: (1) a drop in the average market share of firms exporting product i from origin o to destination d (i.e., ms_{fiobt} goes down) and (2) a rise in the market share of origin o in destination d (i.e., ms_{iodt} goes up). As shown in (6), a drop in ms_{fiobt} increases the demand elasticity, whereas an increase in ms_{iodt} reduces the demand elasticity. So, the overall effect on the demand elasticity and, consequently, the markup can go in either direction in response to a tariff reduction. Whether the elasticity of demand rises or falls will depend on the relative importance of the two channels.

2.2 Markup adjustments and the two reallocation effects

In this subsection, we formalize the idea of two opposing market share effects on markups and discuss how the relative importance of these two effects determines the direction and magnitude of markup adjustments. Our starting point is to decompose changes in markups into two channels: (1) a within-origin reallocation effect which captures the adjustments of markups due to changes in the within-origin market shares, \widehat{ms}_{fiobt} , and (2) an across-origin reallocation effect that captures the markup adjustments due to changes in the across-origin market share, \widehat{ms}_{iodt} .

Proposition 1. *The markup adjustment under our proposed triple-nested CES framework is, up to a first order approximation, given by*

$$\widehat{\mu}_{fiobt} = \underbrace{A(\sigma, \rho, \eta, ms_{fiobt}, ms_{iodt}) \cdot \widehat{ms}_{fiobt}}_{\text{Within-origin reallocation effect}} + \underbrace{B(\sigma, \rho, \eta, ms_{fiobt}, ms_{iodt}) \cdot \widehat{ms}_{iodt}}_{\text{Across-origin reallocation effect}} \quad (11)$$

where the $\widehat{}$ notation represents percentage changes of the variable from one period to the next, i.e., $\widehat{x}_t = \ln(x_{t+1}/x_t)$, and $A(\cdot)$ and $B(\cdot)$ are non-linear scalar functions which take five arguments each (i.e., $\sigma, \rho, \eta, ms_{fiobt}, ms_{iodt}$). Regardless of the initial market share distribu-

tions (i.e., the values of ms_{fiobt} and $ms_{iobt} \forall f, i, o, d$ at t) and the elasticity of substitution across industries η , we have

$$\begin{aligned} A(.) &= B(.) \quad \text{iff} \quad \sigma = \rho \\ A(.) &> B(.) \quad \text{iff} \quad \sigma > \rho \end{aligned} \tag{12}$$

Proof: See appendix B.2.

The key insight of Proposition 1 is that, while $A(.)$ and $B(.)$ are complicated non-linear functions of two market shares (ms_{fiobt}, ms_{iobt}) and all three elasticities (σ, ρ, η), the relative importance of the two market share changes only depends on two elasticities: the within-origin elasticity of substitution σ and the across-origin elasticity of substitution ρ .¹³

When $\sigma = \rho$, we get back to the [Atkeson and Burstein \(2008\)](#) case, where changes in the firm's within-origin market share \widehat{ms}_{fiobt} have exactly the same effect as changes in the origin's market share in the destination \widehat{ms}_{iobt} . In this case, the direction of the markup adjustment depends only on the sum of the two market share changes, i.e., $\widehat{ms}_{fiobt} + \widehat{ms}_{iobt} = \widehat{\omega}_{fiobt}$. This implies that a firm's markup always moves in the same direction as its market share in the destination. Therefore, if a bilateral tariff cut raises the firm's market share in the destination market, it will increase its markup. This is a typical prediction of most oligopolistic competition models but contradicts our empirical findings, which show that a bilateral tariff liberalization can raise firms' market shares but lower their markups.

Our framework allows for a more flexible relationship between the two market shares and markups. When $\sigma > \rho$, the effect of changing the firm's within-origin market share \widehat{ms}_{fiobt} is larger than that of changing the origin's market share in the destination market \widehat{ms}_{iobt} . Therefore, the markup adjustment can be positive even if the sum of the two market share changes is zero or negative (i.e., $\widehat{ms}_{fiobt} + \widehat{ms}_{iobt} = \widehat{\omega}_{fiobt} \leq 0$). Intuitively, this is because when $\sigma > \rho$, firms care a lot more about competition from their peers from the same origin than they do about competition from firms from different origins. Chinese firms, for example, would care a lot about the prices charged by other Chinese firms which export the same product to the destination, but worry much less about the prices charged by competitors from Mexico or Egypt.

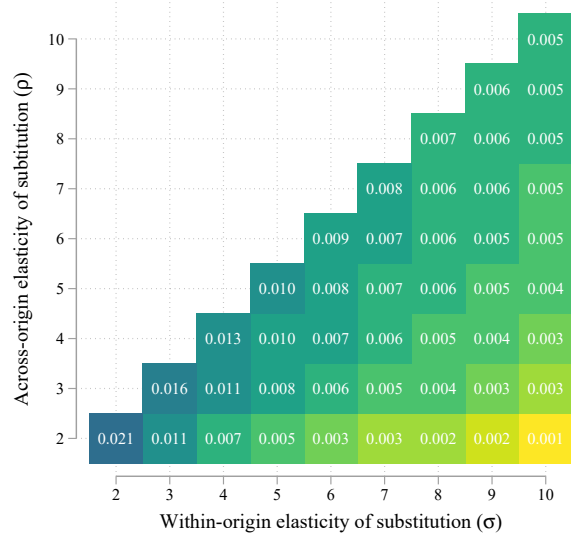
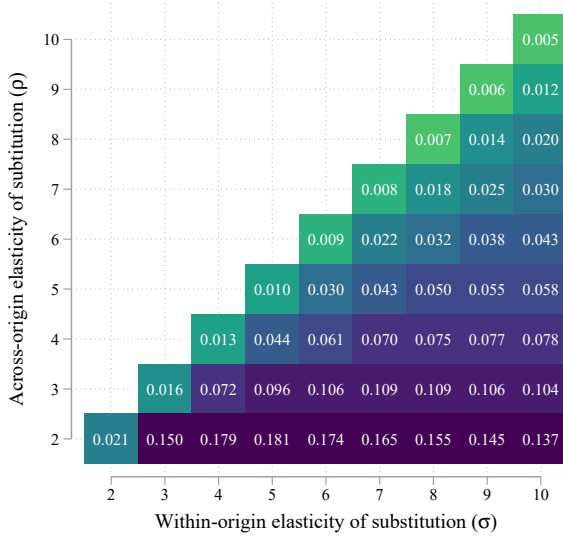
Figure 1 presents a visualization of the $A(.)$ and $B(.)$ functions under different values of within- and across-origin elasticities while fixing the firm's within-origin market share ms_{fiobt} to 50%, the origin's market share in the destination ms_{iobt} to 10% and the elasticity of substitution across products η to 1.2. While the exact quantitative number differs for

¹³Note that this is a general result which holds regardless of underlying shocks that drive the two market share changes.

Figure 1: Visualizing the two reallocation effects on a firm's markup adjustment under different within- and across-origin elasticities (varying ρ and σ while fixing $ms_{fiobt} = 0.5, ms_{iodt} = 0.1$, and $\eta = 1.2$)

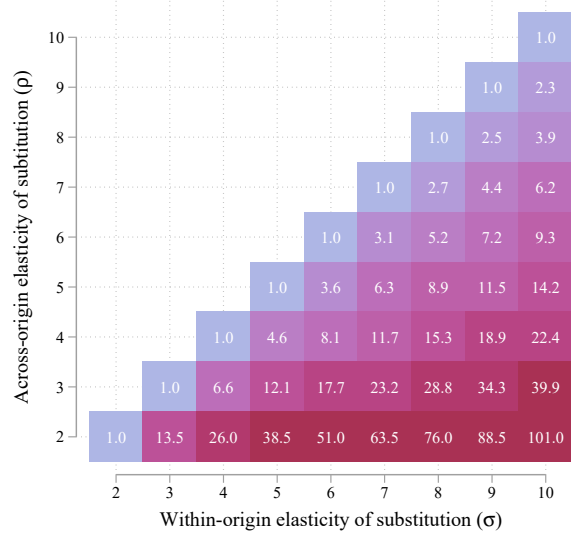
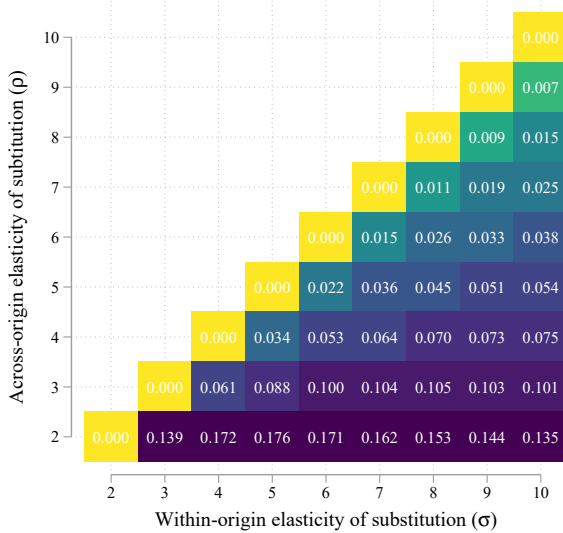
(A) Within-origin reallocation effect
(for a 1% change in the firm's within-origin market share i.e., $\widehat{ms}_{fiobt} = 1\%$)

(B) Across-origin reallocation effect
(for a 1% change in the origin's market share in the destination, i.e., $\widehat{ms}_{iodt} = 1\%$)



(A - B) Level difference of the two effects

(A/B) Ratio of the two effects



Note: The above figures show the values of $A(\sigma, \rho, \eta, ms_{fiobt}, ms_{iodt})$ (top-left), $B(\sigma, \rho, \eta, ms_{fiobt}, ms_{iodt})$ (top-right), $A(.) - B(.)$ (bottom-left) and $A./B(.)$ (bottom-right) varying ρ and σ while fixing $ms_{fiobt} = 0.5, ms_{iodt} = 0.1$, and $\eta = 1.2$. Each colored square indicates the value of the corresponding function (e.g., $A(.)$ for the top-left panel) for a given calibration of the within-origin elasticity of substitution σ and the across-origin elasticity of substitution ρ . The numbers in the coloured cells of the top two figures and the bottom-left figure show the corresponding markup adjustments in percentages. For example, the value 0.021 in the bottom-left cell ($\sigma = 2$ and $\rho = 2$) in the top-left figure reflects a 0.021% markup increase. The numbers in the coloured cells of the bottom-right figure give the ratio of the two reallocation effects (i.e., $A./B(.)$) and are based on a different colour scheme than the other three figures to highlight the different scales.

firms and origins with different market shares, the qualitative pattern remains the same.¹⁴

The top left panel of figure 1 shows percentage changes in the markup for a 1% change in the firm’s within-origin market share, holding the origin’s market share in the destination fixed (i.e., $\widehat{m}s_{fiotd} = 1\%$ and $\widehat{m}s_{iodt} = 0$). Each colored square represents the value of $A(\cdot)$ for a given calibration of the within-origin elasticity of substitution σ , measured on the x-axis, and the across-origin elasticity of substitution ρ , measured on the y-axis. Focusing on the diagonal elements, we can see that the number in the coloured cell goes down as the two elasticities increase, reflecting the fact that firms which sell more substitutable goods have less market power and make smaller markup adjustments for a given change in their market share. While values of the off-diagonal elements (i.e., when $\sigma \neq \rho$) show highly non-linear patterns, we find that the within-origin reallocation effect on markups is in general *more* pronounced the larger the distance between the two elasticities.¹⁵

The top right panel of figure 1 shows percentage changes in the markup for a 1% change in the origin’s market share in the destination, holding firms’ within-origin market shares fixed (i.e., $\widehat{m}s_{fiotd} = 0$ and $\widehat{m}s_{iodt} = 1\%$). As expected, the values of the diagonal elements of $B(\cdot)$ are exactly the same as those of $A(\cdot)$, as the two market share changes have the same effect on markups in the Atkeson and Burstein (2008) model. Intuitively, this is because, when the two elasticities are the same, firms face the same competitive pressure from price adjustments by competitors from their own origin as from other, different origins. Despite the diagonal elements being the same, the off-diagonal elements of $B(\cdot)$ show dramatically different patterns compared to $A(\cdot)$: the across-origin reallocation effect on markups is *less* pronounced the larger the distance between the two elasticities.

The different patterns of the off diagonal elements for functions $A(\cdot)$ and $B(\cdot)$ suggest that the two reallocation effects will not in general cancel out even if the two market shares move in exactly the opposite direction and sum to zero. The bottom left panel of figure 1 shows the percentage change in a markup if the firm’s within-origin market share increases by 1% while the origin’s market share in the destination drops by 1% (i.e., $\widehat{m}s_{fiotd} = 1\%$ and $\widehat{m}s_{iodt} = -1\%$). We can see clearly in the off-diagonal elements that the within-origin reallocation effect dominates when $\sigma > \rho$ as predicted in Proposition 1. Moreover, the magnitude of the level differences of these two effects is largely dictated by the pattern of the within-origin reallocation effect (i.e., $A(\cdot)$) when $\sigma > \rho$.

Finally, the bottom right panel of figure 1 shows the ratio of the two functions ($A(\cdot)/B(\cdot)$), which gives information on the extent to which the origin’s market share in the destination

¹⁴See Appendix B.4 for a discussion of how the market share reallocation effects vary with the firm’s initial within-origin market share and the origin’s initial market share in the destination.

¹⁵Appendix B.3 provides a more detailed discussion of the economic intuition behind these patterns.

\widehat{ms}_{iodt} would need to drop in order to offset the effect of a 1% increase in the firm's within-origin market share \widehat{ms}_{fiodt} on markups. The diagonal elements of 1.0 indicates the origin's market share in the destination \widehat{ms}_{iodt} would need to drop by 1% to offset the effect of a 1% increase in the firm's within-origin market share \widehat{ms}_{fiodt} in the [Atkeson and Burstein \(2008\)](#) case. Focusing on the off-diagonal elements, we can see clearly that the ratio increases dramatically as the distance between the two elasticities becomes larger. At an extreme, when $\rho = 2$ and $\sigma = 10$, the origin's market share in the destination \widehat{ms}_{iodt} would need to drop by more than 100% to offset the effect of a 1% increase in the firm's within-origin market share \widehat{ms}_{fiodt} .

2.3 Sources of within-origin reallocation effects under a bilateral trade policy change

So far, we have argued that the two market share reallocation effects go in opposite directions under a bilateral tariff change. While it is well documented that the origin's market share in the destination ms_{iodt} increases after a bilateral tariff cut, the direction and magnitude of any within-origin market share changes are much less obvious. While the real test of our model mechanism comes from our empirical estimates in section 5, this subsection aims to clarify this question from a theoretical point of view.

The first question we ask is whether we should expect the firm's within-origin market share ms_{fiodt} to change at all. This will help us to understand why the within-origin reallocation effect has been overlooked in the literature. Intuitively, it is tempting to conclude that, since a bilateral tariff change impacts all firms from the origin in the same way, it should have no effect on the relative competitiveness of these firms and thus should not change within-origin market shares.¹⁶ More specifically, we can rewrite the firm's within-origin market share using equations (5) and (7) as:

$$ms_{fiodt} = \frac{\tau_{iodt} \mu_{fiodt} mC_{fiodt} y_{fiodt}}{\sum_{f' \in \mathcal{F}_{iodt}} \tau_{iodt} \mu_{f'iodt} mC_{f'iodt} y_{f'iodt}} = \frac{\mu_{fiodt} mC_{fiodt} y_{fiodt}}{\sum_{f' \in \mathcal{F}_{iodt}} \mu_{f'iodt} mC_{f'iodt} y_{f'iodt}} \quad (13)$$

It is straightforward to see that the bilateral tariff τ_{iodt} drops out from the equation (13), which seems to suggest that a bilateral tariff change does not have first order effects on firms' within-origin market shares. Therefore, for a bilateral tariff change, the within-origin reallocation effect should be approximately zero.

¹⁶For example, [Amiti, Itskhoki and Konings \(2019\)](#) note that firms pass-through common and idiosyncratic shocks very differently in a [Atkeson and Burstein \(2008\)](#) model: a common cost shock has a limited impact on the market structure (i.e., the market share distributions) and thus leads to full pass through of the shock with limited markup adjustments, while an idiosyncratic cost shock leads to significant changes in market shares and therefore large markup adjustments.

While the above argument is approximately accurate in a model with a fixed set of firms before and after the tariff change ($\mathcal{F}_{iodt} = \mathcal{F}_{iodt+1}$), it fails to consider an important aspect of international trade – firms’ endogenous market participation choices – which is a key margin along which firms respond to bilateral trade policy changes. If a tariff cut encourages entry, then the denominator of (13) will increase non-trivially, reducing the market shares of existing exporters in the destination. Facing more competitive pressure from their peers, the lower within-origin market shares of existing firms translate to a higher demand elasticity and a lower desired markup.

A natural follow-up question is whether the entry and exit of relatively small firms would lead to economically meaningful changes in the within-origin market structure (and the denominator of (13)). Intuitively, in a trade model with a continuous measure of heterogeneous firms like Melitz (2003), entry and exit of marginal firms would have a negligible direct impact on the existing exporters’ market shares. This may lead some to conclude that the within-origin reallocation effect must be very small even if it is not zero.

A key departure of our model from existing trade models, as discussed in subsection 2.1, is that we envision a very small number of firms at the highly disaggregated product-origin-destination level. Our design is motivated by our data, which shows that the median number of exporters that operate at the product-origin-destination level is two (see Appendix table A8 for more details). Due to the granularity and discreteness of the firm distribution, the entry and exit of an exporter can have very big impacts on the within-origin market structure. For example, for a destination with two similar firms from the an origin selling the same product, the entry of a third firm as a result of a bilateral tariff cut could cause each of the existing firms to lose up to 33% [= $(0.5 - 0.33)/0.5$] of their initial within-origin market share.

2.4 Welfare implications

A key takeaway from subsections 2.2 and 2.3 is that exporters’ markup adjustments in response to a bilateral tariff cut can be either positive or negative in our more general framework as a result of the two opposing reallocation effects. In particular, incumbent exporters will tend to reduce their markups in response to a bilateral tariff cut when (i) the varieties produced by an origin are more substitutable with each other than with those produced by different origins ($\sigma > \rho$) and (ii) the number of exporters selling the same product from the origin in the destination market is small.

This finding has important implications for pro-competitive gains from trade. As recently pointed out by Arkolakis, Costinot, Donaldson and Rodríguez-Clare (2018), the pro-

competitive effects of trade depend crucially on how exporters and domestic firms adjust their markups. They show that, in a large set of trade models, the pro-competitive gains from trade are elusive, as the gains from domestic firms' markup reductions are offset by the losses due to foreign exporters' markup increases. In our conceptual framework with a discrete number of firms and entry and exit, we find that exporters may not increase their markups by much after a bilateral tariff cut, as a result of increased competitive pressure from new entrants. In fact, under a large set of calibrations discussed in section 2.2, incumbent exporters tend to *lower* their markups in response to a bilateral tariff cut. This implies that there are no effects which offset the pro-competitive gains from trade – these gains are strictly positive since both incumbent exporters and surviving domestic firms reduce their markups.

While our theoretical framework highlights the possibility that both reallocation effects are active, whether the within-origin reallocation effect is strong enough to dominate the overall direction of markup adjustments remains an empirical question.

3 Data

We bring together information on firms' product-level export values and quantities for eleven origins, 83 preferential trade agreements, and bilateral tariffs for 165 destinations to estimate the effect of trade policy on firms' exporting behaviour and markups. Our final dataset contains 25,176,098 observations at the firm-product-origin-destination-year level and spans the years 2000-2013.

3.1 Firm-level trade

We use administrative data on the universe of product exports by firms for eleven developing and emerging economies, obtained from three different sources. Data for Albania, Bulgaria, Burkina Faso, Malawi, Mexico, Peru, Senegal, Uruguay and Yemen are taken from the World Bank Exporter Dynamics Database, data for Egypt from the Economic Research Forum Exports Dataset and data for China from the Chinese Customs Database.¹⁷ While data for different countries are available for different years, as summarised in table 1, 88% of observations in our final dataset are from 2000-2006.

Apart from the Chinese Customs Database, which contains monthly data for HS8 products, the raw datasets provide information on non-zero annual firm level export values and

¹⁷For more information about the World Bank Exporter Dynamics Database, see [Cebeci, Fernandes, Freund and Pierola \(2012\)](#) and [Bortoluzzi, Fernandes and Pierola \(2015\)](#).

volumes to individual foreign destinations by HS6 product. Export values are provided in US dollars and reported on a FOB basis for all countries except Senegal, which reports CIF figures. Export volumes represent net weight in kilograms, with the exception of China and Egypt, which use a variety of measures, as well as Mexico, which does not specify the measures used between 2000 and 2009. To ensure that our data are comparable across our eleven origin countries, we aggregate the monthly Chinese data to the annual level. For all eleven countries, we drop observations for which we cannot determine the destination country, observations which report a product code that is not part of any HS revision during our sample period and observations with missing or negative reported trade values.¹⁸ As our dataset spans multiple revisions of the HS classification system, we further convert the raw HS6 codes to consolidated HS codes which are stable over time (see appendix A.2 for more details). Similar to other studies using administrative data, we use trade unit values as a proxy for prices.¹⁹

3.2 Trade policy

We source data on trade agreements from the World Bank Deep Trade Agreements (WB DTA) Database and data on preferential and most favoured nation (MFN) tariffs from the WTO Integrated Database (WTO IDB). To capture the phase-in of trade agreements, we supplement the data sourced from the WTO IDB with information contained in the tariff data compiled by [Feenstra and Romalis \(2014\)](#).

The WB DTA contains detailed information on the contents of trade agreements, their members and the years they were adopted as well as, where applicable, discarded, for 257 agreements which entered into force between 1958 and 2015. The eleven countries in our sample are involved in 83 of these trade agreements, 28 of which entered into force during our sample period. We use information on these agreements to construct an indicator variable that records whether there is an active trade agreement between an origin and a destination in our sample in any given year.

The WTO IDB contains HS6-product-level data on preferential and applied MFN ad-valorem tariffs for the years 2000-2013 for 138 and 165 destination countries, respectively. We aggregate the raw data to consolidated HS codes by taking a simple average across HS6 codes. To address missing values, we follow [Feenstra and Romalis \(2014\)](#). For applied MFN tariffs, we replace missing values with the closest preceding value, on the basis that updated

¹⁸Additionally, we drop exports from China to Hong Kong, which likely acts as an entrepot during this period.

¹⁹To address any issues that might arise if different quantity measures were reported in different datasets, we include firm-product-origin-time fixed effects in our markup regressions (see section 4).

Table 1: Firm-level trade data: countries and years

Country	Years	Firms	Observations	... with PTA	... with Tariff
Albania	2004 - 2012	6,314	66,397	6,090	65,330
Bulgaria	2001 - 2006	50,780	780,816	99,789	733,662
Burkina Faso	2005 - 2007	718	6,492	3,413	6,307
	2008 - 2012	1,173	10,305	6,016	10,113
China	2000 - 2006	230,339	20,043,162	1,168,391	19,221,840
Egypt	2005 - 2013	20,461	612,907	496,316	494,124
Malawi	2006 - 2008	1,360	9,409	5,903	9,104
	2009 - 2012	3,036	20,536	13,818	20,107
Mexico	2000 - 2007	106,688	1,904,144	1,230,160	1,861,198
	2008 - 2009	44,971	635,065	399,090	625,321
	2010 - 2011	43,866	678,719	415,385	667,921
	2012	32,706	390,582	308,744	385,475
Peru	2000 - 2013	28,851	888,886	339,287	854,547
Senegal	2000 - 2012	2,919	82,275	44,955	80,698
Uruguay	2001 - 2012	7,300	132,844	45,210	124,256
Yemen	2008 - 2012	1,242	18,850	11,533	16,095

Notes: The datasets for Burkina Faso, Malawi and Mexico feature multiple distinct panels as a result of changes to the system of firm identifiers. The columns “...with PTA” and “...with Tariff” report the number of observations for which our binary PTA_{odt} variable takes a positive value and our $Tariff_{iodt}$ variable takes non-missing values. For PTA_{odt} , this amounts to the number of observations for which there is an active PTA between the origin and the destination. For $Tariff_{iodt}$, this refers to the number of observations for which data on bilateral tariffs is available.

tariff schedules are more likely to be available after significant changes. In cases where there is no preceding value, we use the closest subsequent value. For preferential tariffs, we extract information about the phase-in of trade agreements from the dataset compiled by [Feenstra and Romalis \(2014\)](#), and then use this data to impute missing values (see appendix [A.1](#) for details). We then set our bilateral tariff variable equal to the lowest reported preferential tariff a destination offers to exporters from a given origin, when it is available, and use data on the MFN tariff applied by the destination, when it is not. The eleven national customs databases report exports to a total of 251 foreign destinations. Omitting observations for the smaller destinations for which no tariff data is available reduces the size of the final estimation dataset from 26,281,389 to 25,176,098.

4 Empirical strategy

We identify elasticities to trade policy changes by estimating the following reduced form specification:

$$\ln(\text{Outcome}_{fiodt}) = \beta_1 \cdot \text{PTA}_{odt} + \beta_2 \cdot \text{Tariff}_{iodt} + \delta_{fiot} + \delta_{idt} + \delta_{od} + \epsilon_{fiodt}, \quad (14)$$

where the outcome variables vary over five dimensions – firm f , product i , origin country o , destination country d , and year t .²⁰ The two right hand side variables describe the trade policy regime firms from an origin face in the destination. The first, PTA_{odt} , is an indicator variable equal to one if the origin and the destination have an active trade agreement in year t . The second, Tariff_{iodt} , denotes the natural logarithm of one plus the ad-valorem tariff on imports of product i from origin o charged by destination d .

We study the responsiveness of key outcome variables suggested by our model. The Atkeson-Burstein model of international pricing emphasizes changes in a firm’s market share in a destination, i.e., $\omega_{fiodt} = p_{fiodt}y_{fiodt} / \sum_{f,o} p_{fiodt}y_{fiodt}$ where the summation in the denominator includes domestic sales of firms in the destination country. To estimate the elasticity of a firm’s market share in a destination to trade policy, we rearrange the definition of market share, take logs, and note that we can control for unobservable product-level domestic

²⁰All continuous outcome variables enter our estimating equations in natural logarithms. We estimate these specifications with panel OLS rather than PPML regressions, as creating a full panel of zero trade flows at the firm-product-origin-destination-year level for eleven countries would result in a dataset of several billion observations, rendering estimation infeasible in a reasonable time period. This means that we only use observations with positive trade flows. However, it is important to note that our fixed effects should absorb most of the variation in trade costs that prevents firms from entering a market, and thus account for the selection process that gives rise to the positive trade flows, as noted in [Baier, Bergstrand and Feng \(2014\)](#) and [Corsetti, Crowley, Han and Song \(2021\)](#).

consumption in the destination with product-destination-time fixed effects:

$$\ln(p_{fiodt}y_{fiodt}) = \ln(\omega_{fiodt}) + \underbrace{\ln\left(\sum_{f,o} p_{fiodt}y_{fiodt}\right)}_{\text{absorbed by } idt \text{ fixed effects}} \quad (15)$$

Thus, when the outcome is the log of (tariff inclusive) export sales, $\ln(p_{fiodt}y_{fiodt})$, and firm-product-origin-time (δ_{fiot}), product-destination-time (δ_{idt}) and origin-destination (δ_{od}) fixed effects are included in specification (14), β_1 and β_2 can be interpreted as the elasticity of a firm's market share in a destination to PTA participation and the tariff, respectively.²¹ The inclusion of firm-product-origin-time (δ_{fiot}) and product-destination-time (δ_{idt}) fixed effects means that the firm's market share elasticities to trade policy are identified from cross-sectional variation in a firm's market share for a product across different destinations as well as variation across different firms and origins.²² The origin-destination fixed effect δ_{od} , meanwhile, absorbs any variation due to the distance between two countries as well as their geography, history and culture.²³

The next variable of interest is the firm's markup. To estimate the elasticity of firms' markups to trade policy, we decompose their overall markups into FOB prices and marginal costs. As before, we can rearrange this definition and control for the latter with firm-product-origin-time fixed effects:

$$\ln(p_{fiodt}^b) = \ln(\mu_{fiodt}) + \underbrace{\ln(mc_{fiot})}_{\text{absorbed by } fiot \text{ fixed effects}} \quad (16)$$

When the outcome variable is the natural logarithm of the FOB unit value excluding tariffs and other trade costs, $\ln(p_{fiodt}^b)$, and firm-product-origin-time (δ_{fiot}), product-destination-time (δ_{idt}), and origin-destination (δ_{od}) fixed effects are used in equation (14), the β parameters therefore identify the elasticity of the firm's markup for a product in a destination to trade policy changes in that destination. This is because the firm-product-origin-time fixed effect δ_{fiot} controls for a firm's time-varying productivity, marginal costs and global markup common to all foreign destinations, while the product-destination-time

²¹We use tariff inclusive export sales because both the [Atkeson and Burstein \(2008\)](#) model and our theoretical framework define the firm's market share in terms of the price paid by consumers, which includes the tariff.

²²The product-origin-time element of the firm-product-origin-time fixed effects δ_{fiot} and the product-destination-time fixed effects δ_{idt} are standard tools to capture multilateral resistance terms in the gravity literature ([Anderson and van Wincoop, 2003](#); [Feenstra, 2004](#); [Redding and Venables, 2004](#); [Head and Mayer, 2014](#); [Baier, Bergstrand and Feng, 2014](#)).

²³Origin-destination fixed effects δ_{od} also absorb pricing variation associated with time-invariant features such as quality ([Bastos and Silva \(2010\)](#)) or, for instance, the Alchian-Allen effect ([Hummels and Skiba \(2004\)](#)).

fixed effect δ_{idt} absorbs time-varying aspects in the local destination such as market size, the industry price level (p_{idt}) and demand shifts (y_{idt}).

The final outcome variable we examine with (14) is the natural logarithm of a firm's *share of trade* in a given product between its origin and the destination, $\ln(ms_{fiot})$. This variable captures one of the two reallocation effects suggested by the triple-nested CES preference structure we introduce in this paper. A positive sign on the tariff coefficient would imply that, on average, firms' share of their origin's trade rises with the tariff, and support the idea that competition among firms from the same origin tends to be less intense and more oligopolistic when tariffs are higher. In this specification, the inclusion of firm-product-origin-time fixed effects δ_{fiot} implies that elasticities are identified from cross-destination variation in a firm's product-level share of its origin's trade with the destination, while product-destination-time fixed effects δ_{idt} again control for market-specific shifts in demand. Finally, as before, an origin-destination fixed effect δ_{od} absorbs time-invariant features that influence bilateral trade.

The second reallocation effect variable suggested by our triple-nested CES preference structure, the natural logarithm of the origin's market share in the destination, $\ln(ms_{iodt})$, requires a reduced form specification at a different level of aggregation, i.e.,

$$\ln(\text{Outcome}_{iodt}) = \beta_1 \cdot \text{PTA}_{odt} + \beta_2 \cdot \text{Tariff}_{iodt} + \delta_{iot} + \delta_{idt} + \delta_{od} + \epsilon_{iodt} \quad (17)$$

The elasticity of the origin's market share in a destination with respect to trade policy is identified using a strategy similar to that employed to estimate the firm's market share in a destination. We begin with the definition of an origin's market share for a product i in a destination, $ms_{iodt} = p_{iodt}y_{iodt} / \sum_o p_{iodt}y_{iodt}$, where the summation in the denominator includes domestic producers. The product-destination-time fixed effect δ_{idt} included in specification (17) controls for total consumption of product i in destination d so that using the country's tariff inclusive exports $\ln(p_{iodt}y_{iodt})$ as the outcome variable in equation (17) identifies β_1 and β_2 as the elasticities of the origin's market share in the destination, ms_{iodt} , to trade policies. Note that shifts in supply or demand in the origin country are absorbed through the inclusion of product-origin-year (δ_{iot}) and product-destination-year (δ_{idt}) fixed effects and that origin-destination fixed effects δ_{od} control for gravity forces.

The final piece of the puzzle in our analysis is an examination of firm entry and exit - summarized by the number of firms from the origin selling a given product to the destination, $N_{iodt} \equiv |\mathcal{F}_{iodt}|$. For this variable, we create a dataset at the product-origin-destination-year level which includes zeros in years when no firms are observed exporting a product that the origin exports to the destination in at least one year and use PPML to estimate specification

(17).

A general observation on the use of fixed effects in the specifications discussed above, and particularly the origin-destination fixed effect δ_{od} , is that they have an additional benefit of addressing potential endogeneity problems. For example, the existence of a PTA could be intertwined with the level of bilateral trade flows – the larger the trade flows between two countries, the greater the benefits from and therefore the incentive to sign a PTA. This means that there is potential for reverse causality, and that it might be large trade flows which cause PTAs, rather than PTAs which cause large trade flows. While this is unlikely to be a problem at the firm level, it could be an issue at the country-pair level. Accounting for unobserved heterogeneity at this level should therefore all but resolve these concerns (see Baier and Bergstrand (2007)).

5 Empirical results

We find that exporting firms respond to the tariff liberalizations associated with preferential trade agreements by lowering their markups. The richness of our multi-origin panel allows us to trace out not only changes in markups, but also the role of the two different market share measures that influence the elasticity of demand facing a firm under our triple-nested preference structure. We show that preferential tariff liberalizations stimulate entry from an origin to such a degree that the market power of individual firms from that origin declines in the destination, even as the total market share of the origin in the destination rises. This is an exciting result which highlights the importance of examining precisely how oligopolistic competition evolves under a trade liberalization and which could help explain the puzzling empirical finding that the tariff-exclusive prices of Chinese exports were virtually unchanged in the face of US tariffs imposed as part of the US-China Trade War.²⁴

²⁴Although the US-China Trade War studies of both [Amiti, Redding and Weinstein \(2019\)](#) and [Fajgelbaum, Goldberg, Kennedy and Khandelwal \(2020\)](#) found large declines in the value of trade from China, their analyses of import unit values from China showed almost no decline in response to steep tariff hikes. This type of phenomenon could arise if the two reallocation effects discussed in subsection 2.2 have offsetting impacts on prices. That is, on the one hand, the US tariff hike can induce less productive Chinese producers to exit the US market. Therefore, continuing Chinese producers would face less competitive pressure from their Chinese peers, and their markups would tend to rise. This is the within-origin reallocation effect. On the other hand, continuing Chinese producers would become less competitive relative to firms from competing origins due to the direct effect of the tariff hike on Chinese merchandise, which would lead them to lower their markups. This is the across-origin reallocation effect. In our empirical analysis, for which Chinese data is available over 2000-2006, we find the within-origin reallocation force dominates the overall effect on markups. Over 2000-2006, it may have been that the varieties made by Chinese producers were substantially more substitutable among themselves than with the products made by other countries. As China's economy shifted towards private enterprises and more firms started exporting, the number of Chinese firms in any given destination's product market would have increased, so that by 2018 any exit associated with tariff changes would be expected to have only a small impact on the incumbent firms' markups. Under this

Table 2: Elasticities of market shares, markups, and counts of firms to tariffs

	Firm's mkt share in the dest. $\ln(\omega_{fiotd})$ (1)	Markups $\ln(\mu_{fiotd})$ (2)	Firm's within origin mkt share $\ln(ms_{fiotd})$ (3)	Origin's mkt share $\ln(ms_{ioidt})$ (4)	No. of firms (PPML) (5)
Tariff _{ioidt}	-0.78*** (0.242)	0.41*** (0.073)	2.86*** (0.322)	-2.31*** (0.271)	-2.21*** (0.162)
PTA _{oidt}	0.01 (0.022)	-0.02** (0.008)	0.05* (0.027)	-0.04* (0.023)	-0.05*** (0.009)
R ²	0.66	0.90	0.79	0.77	-
Observations	15,853,618	15,793,386	15,853,618	1,067,240	2,750,833
Fixed Effects					
Firm-product-origin-year	✓	✓	✓		
Product-origin-year				✓	✓
Product-destination-year	✓	✓	✓	✓	✓
Origin-destination	✓	✓	✓	✓	✓

Notes: The dependent variable is the firm's log (tariff-inclusive) export value in column (1), the firm's log (tariff-exclusive) unit value in column (2), the log of the firm's share of its country's trade with the destination in column (3), the log of the country's (tariff-inclusive) export value to the destination market in column (4) and the number of firms in column (5). Tariff_{ioidt} and PTA_{oidt} capture the trade policy firms from the origin face in the destination. Standard errors, reported in parentheses, are clustered at the product-destination level, and we denote statistical significance with *** p<0.01, ** p<0.05, and * p<0.1. Estimates are based on an integrated dataset of firms' exports from eleven countries built from the World Bank Exporter Dynamics Database, China's Customs Authority, and Egypt's Customs Authority, as well as tariff data from the WTO and [Feenstra and Romalis \(2014\)](#), and the World Bank Deep Trade Agreements Database.

We present our main results, estimates of elasticities to bilateral tariffs and PTA participation, in table 2. The first column contains the elasticity of a firm’s market share in a destination, at the level of an HS6 product, to the tariff it faces.²⁵ A 10% reduction in tariffs is associated with an 8% increase in a firm’s market share in the destination. For a firm with an initial market share of 50%, this means that a 10% tariff cut, for example from 10% to 9%, will increase the firm’s market share to 54%. This shows that the bilateral tariff cut increases the market power of firms from the preferred origin at the expense of firms exporting from other origins as well as domestic firms. Recall that in the [Atkeson and Burstein \(2008\)](#) model of oligopolistic competition, which emphasizes competition among all firms within a product market, the effect of tariffs on firms’ market shares would be a sufficient statistic for the direction of the change in markups in response to a change in trade policy. Notably, a tariff cut would imply an increase in the firm’s market share in the destination, a decline in the demand elasticity it faces, and a rise in the firm’s markup. This is not what we find empirically. Turning to markups, column (2) identifies how the component of the markup that is specific to a destination, the residual component of the markup that varies across destination markets, changes when a country joins a PTA.²⁶ Signing a PTA reduces markups by 2%. More interestingly, we find that a 10% reduction in the tariff on a product is associated with a 4% *decline* in firms’ markups. This finding, which shows that markups fall while market shares rise, directly contradicts the prediction of most standard models of oligopoly.²⁷

Our theoretical model highlights the importance of decomposing the firm’s market share in the destination into two parts, the origin country’s share of the destination market (ms_{iodt}) and the firm’s share of its country’s trade with the destination (ms_{fiomt}). In a world with oligopolistic competition that is shaped by the substitutability of varieties both across origins within an industry and across firms within an origin, markups depend on changes in both of these market shares. We see the negative sign on the elasticity of firms’ overall market share in the destination to tariffs (column 1) can be decomposed into a positive sign on the

scenario, by the time of the US-China Trade War, the within-origin reallocation effect would have become weaker so that the two reallocation effects could have just offset each other to result in the zero price change documented in the recent literature.

²⁵As discussed in section 4, regressing the log value of a firm’s product- and destination-specific exports on a product-destination-year fixed effect implies that the parameter estimate on the tariff captures the elasticity of a firm’s market share to the tariff (see equation (15)).

²⁶The inclusion of firm-product-origin-time fixed effects controls for time-varying marginal costs at the level of the product within a firm as well as time-variation in the global or common markup that the firm charges in all foreign destinations (see section 4).

²⁷As discussed in [Helpman and Krugman \(1985\)](#), the results in trade models of oligopoly are extremely sensitive to precise market structures. That said, most standard quantity competition models yield a positive correlation between market shares and markups.

elasticity of the firm’s share of its country’s trade with the destination (column 3) and a negative sign on the elasticity of the origin’s share of the destination market (column 4). Using a traditional definition of import market share, that of the origin in the destination (see column 4), we find a country’s import market share rises 23% when it is the beneficiary of a 10% preferential tariff cut. As first described by Viner (1950), when one country enjoys a tariff cut in a destination that is not offered to competing origins, the country’s market share in that destination rises.

But building on this with firm-level data, our analysis shows more subtle forces are at play. Turning to a firm’s trade share in a destination among all firms from its own origin, we find that a 10% reduction in the bilateral tariff is associated with a substantial *decline* in the average market share of an exporting firm of 29% (see column (3)). Importantly, market participation by exporting firms increases as the bilateral tariff falls (column (5)). A 10% cut in the tariff imposed by a destination leads to a 22% increase in the number of exporters from the affected origin. The strong extensive margin response from the origin affects both market shares that our theoretical model suggests influence the impact of a trade liberalization on markups, an origin’s share of trade in a destination and a firm’s share of trade among compatriot firms from its origin, and moves them in opposite directions.

Interpreted through the lens of our triple-nested CES model, this suggests that a tariff liberalization leads to an across-origin reallocation effect that puts upward pressure on markups and a within-origin reallocation effect that puts downward pressure on markups. As discussed in section 2, the net effect on markups will depend on which of these two effects dominates. The finding that markups decline with tariff cuts implies that the within-origin reallocation effect dominates the across-origin reallocation effect, and that the elasticity of demand facing a firm therefore falls, in our sample. This is consistent with the idea that consumers’ preferences across varieties lead firms in our dataset to view firms from their own origin as more relevant competitors in the destination market than firms from other origins and react more strongly to additional entrants from their own origin than to the fall in their trade costs in setting prices.

Table 2 also shows that preferential trade agreements have small effects on some of our outcomes of interest beyond the tariff reductions they embody. PTAs signed by low and middle income countries typically involve much larger tariff cuts than those among high income countries, suggesting that most of the benefit of a preferential trade agreement for low and middle income countries comes from tariff cuts rather than provisions that simplify cross-border trade or remove non-tariff barriers.²⁸ It is therefore not surprising that the direct effect of the PTA dummy in our dataset is limited. Finally, we present various robustness

²⁸We thank Jeff Bergstrand for sharing this insight.

checks for our baseline results in appendices C.2 and C.3.

5.1 Trade policy and product differentiation

To investigate whether firms' and markup responses vary systematically with the degree of substitutability of a product, as predicted by our theoretical framework and documented in prior work (Corsetti, Crowley, Han and Song (2021)), we use the CCHS commodity classification system to split our sample into highly and less differentiated products.²⁹ Our framework predicts that firms which sell highly differentiated goods, and therefore operate in markets in which there is considerable scope to exploit market power, should adjust markups more than firms which sell less differentiated products that are highly substitutable.

We explore this idea in table 3, which reports exporters' responses to changes in trade policy by degree of product differentiation. The top panel presents results for the subsample of highly differentiated goods and the bottom panel presents results for the subsample of goods which are less differentiated, including commodities and simple manufactured goods like processed foods. We begin by discussing the elasticities of the firm's within-origin market share and the origin's market share to the tariff reported in columns (3) and (4). We consistently find a positive sign on the elasticity of the firm's within-origin market share and a negative one on the origin's market share for both highly and less differentiated goods. Although the magnitudes of the elasticities are different for highly and less differentiated goods, within each type of good the magnitudes of the two elasticities are similar: for a 10% bilateral tariff cut, the firm's within-origin market share drops by 36.4% (15.3%) while the origin's market share increases by 34.5% (15.6%) for highly (less) differentiated goods. If we were to look at the firm's within-origin market share change and the origin's market share change through the lens of the Atkeson and Burstein (2008) model, we would predict markup adjustments of zero as the two reallocation effects would cancel out (see the diagonal elements in the bottom left panel of figure 1).

However, empirically, we find that markups adjust substantially for highly differentiated goods but do not for less differentiated goods (see column (2)). This finding is consistent with two predictions of our more general model discussed in section 2. First, our model predicts

²⁹ Most studies adopt the industry classifications set forth by Rauch (1999), according to which a product is differentiated if it does not trade on organized exchanges and/or its price is not regularly published in industry sales catalogues. While this system is quite powerful in identifying commodities, a drawback is that the vast majority of manufactured goods end up being classified as differentiated. The CCHS classification refines the class of differentiated goods in Rauch into two categories – highly and less differentiated. Corsetti, Crowley, Han and Song (2021) calculate that in the Chinese Customs Database 2000-2014, 79.8 percent of observations are classified by Rauch as differentiated. Of these, only 48.6 percent are categorized as highly differentiated under the CCHS Chinese-linguistics-based classification system. See appendix A.3 for further details.

Table 3: Highly vs. less differentiated goods

	Firm's mkt share in the dest. $\ln(\omega_{fi\text{odt}})$ (1)	Markups $\ln(\mu_{fi\text{odt}})$ (2)	Firm's within origin mkt share $\ln(ms_{fi\text{odt}})$ (3)	Origin's mkt share $\ln(ms_{i\text{odt}})$ (4)	No. of firms (PPML) (5)
Highly Differentiated Goods					
Tariff _{iodt}	-1.39*** (0.386)	0.87*** (0.106)	3.64*** (0.430)	-3.45*** (0.363)	-2.80*** (0.197)
PTA _{odt}	-0.00 (0.041)	-0.03* (0.016)	0.20*** (0.053)	-0.20*** (0.038)	-0.04** (0.017)
R ²	0.62	0.93	0.75	0.80	-
Observations	5,803,447	5,792,021	5,803,447	346,253	891,704
Less Differentiated Goods					
Tariff _{iodt}	-0.07 (0.287)	0.06 (0.077)	1.53*** (0.387)	-1.56*** (0.383)	-1.43*** (0.184)
PTA _{odt}	0.07*** (0.027)	-0.03*** (0.011)	-0.05 (0.031)	0.04 (0.030)	0.02** (0.011)
R ²	0.70	0.90	0.77	0.75	-
Observations	7,800,002	7,758,623	7,800,002	677,634	1,771,815
Fixed Effects					
Firm-product-origin-year	✓	✓	✓		
Product-origin-year				✓	✓
Product-destination-year	✓	✓	✓	✓	✓
Origin-destination	✓	✓	✓	✓	✓

Notes: The dependent variable is the firm's log (tariff-inclusive) export value in column (1), the firm's log (tariff-exclusive) unit value in column (2), the log of the firm's share of its country's trade with the destination in column (3), the log of the country's (tariff-inclusive) export value to the destination market in column (4) and the number of firms in column (5). Products are separated into highly differentiated and less differentiated goods based on the CCHS classification system. Standard errors, reported in parentheses, are clustered at the product-destination level, and we denote statistical significance with *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$. Estimates are based on an integrated dataset of firms' exports from eleven countries built from the World Bank Exporter Dynamics Database, China's Customs Authority, and Egypt's Customs Authority, as well as tariff data from the WTO and [Feenstra and Romalis \(2014\)](#), and the World Bank Deep Trade Agreements Database.

that the within-origin reallocation effect tends to be larger and is more likely to dominate the markup adjustment for goods whose σ and ρ have lower values. This is particularly relevant when the two market share changes are of similar magnitude. As illustrated in the bottom left panel of figure 1 (A-B), the net impact on markups coming from the two reallocation effects are at their largest when the across-origin elasticity of substitution ρ is around 2 and the within-origin elasticity of substitution σ is around 4 or 5. Second, since the two market share effects do not in general cancel each other out, our model predicts that the larger the magnitudes of the two market share changes, the bigger the markup adjustment. Because we find empirically that market share adjustments for highly differentiated goods are more than twice the size of those for the less differentiated goods, we expect markups of highly differentiated goods to change more.

5.2 Pro-competitive trade agreements and global value chains

In this section, we introduce a new dimension to refine our breakdown of the product space and explore the role of PTA participation in global value chains. We do this by using the Broad Economic Categories classification (Rev. 4) to distinguish between intermediate inputs and final consumption goods. Firms which produce and sell final consumption goods often engage in activities such as marketing or branding that aim to differentiate their product relative to their competitors in the marketplace. This suggests that markets for final consumption goods might be oligopolistic with firms that operate in these markets holding some amount of market power. As a result, we would expect changes in barriers to entry to have large impacts on market shares and markups. In contrast, intermediate goods might include more classes of products that are similar to commodities or be more substitutable across varieties.

Table 4 presents results for consumption goods, and table 5 for intermediate inputs. They each contain three panels: the top panel considers all consumption or intermediate goods, the middle panel hones in on highly differentiated goods and the bottom panel reports results for less differentiated goods. Comparing the top panels of the two tables, we see that a 10% tariff reduction increases markups for consumption goods by 6%, but has no effect on the markups of intermediate inputs. We also see that the three different market shares, and particularly firms' share of their origin's trade, as well as the number of firms in the market, respond more strongly for consumption goods. Turning to the middle two panels, the effects on consumption goods, but not on intermediates, appear to be almost entirely driven by highly differentiated consumption goods, for which a 10% tariff reduction leads to a 10% increase in markups. The effect of a tariff liberalization on market shares and market

Table 4: Trade policy elasticities and global value chains: Consumption goods

	Firm's mkt share in the dest. $\ln(\omega_{fiodt})$ (1)	Markups $\ln(\mu_{fiodt})$ (2)	Firm's within origin mkt share $\ln(ms_{fiodt})$ (3)	Origin's mkt share $\ln(ms_{iodt})$ (4)	No. of firms (PPML) (5)
Consumption Goods					
Tariff _{<i>iodt</i>}	-0.88*** (0.335)	0.63*** (0.091)	4.33*** (0.385)	-1.96*** (0.400)	-3.08*** (0.219)
PTA _{<i>odt</i>}	0.12*** (0.038)	-0.03** (0.012)	0.26*** (0.046)	-0.11*** (0.040)	-0.07*** (0.017)
R ²	0.65	0.92	0.75	0.79	-
Observations	6,876,997	6,872,965	6,876,997	377,696	927,593
Highly Differentiated Consumption Goods					
Tariff _{<i>iodt</i>}	-1.15** (0.476)	1.00*** (0.129)	5.01*** (0.534)	-2.47*** (0.428)	-3.18*** (0.224)
PTA _{<i>odt</i>}	0.04 (0.052)	-0.05*** (0.017)	0.40*** (0.073)	-0.30*** (0.052)	-0.15*** (0.025)
R ²	0.60	0.92	0.74	0.82	-
Observations	4,075,000	4,074,107	4,075,000	186,673	450,586
Less Differentiated Consumption Goods					
Tariff _{<i>iodt</i>}	-0.45 (0.496)	0.11 (0.122)	1.41*** (0.465)	-0.47 (0.665)	-1.54*** (0.325)
PTA _{<i>odt</i>}	0.21*** (0.051)	-0.02 (0.016)	0.14*** (0.053)	0.10* (0.060)	0.01 (0.021)
R ²	0.70	0.92	0.78	0.78	-
Observations	2,643,838	2,640,776	2,643,838	185,495	463,386
Fixed Effects					
Firm-product-origin-year	✓	✓	✓		
Product-origin-year				✓	✓
Product-destination-year	✓	✓	✓	✓	✓
Origin-destination	✓	✓	✓	✓	✓

Notes: The dependent variable is the firm's log (tariff-inclusive) export value in column (1), the firm's log (tariff-exclusive) unit value in column (2), the log of the firm's share of its country's trade with the destination in column (3), the log of the country's (tariff-inclusive) export value to the destination market in column (4) and the number of firms in column (5). Products are separated into different groups based on the CCHS and BEC (Revision 4) classification systems. Standard errors, reported in parentheses, are clustered at the product-destination level, and we denote statistical significance with *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$. Estimates are based on an integrated dataset of firms' exports from eleven countries built from the World Bank Exporter Dynamics Database, China's Customs Authority, and Egypt's Customs Authority, as well as tariff data from the WTO and [Feenstra and Romalis \(2014\)](#), and the World Bank Deep Trade Agreements Database.

Table 5: Trade policy elasticities and global value chains: Intermediate inputs

	Firm's mkt share in the dest. $\ln(\omega_{fiotd})$ (1)	Markups $\ln(\mu_{fiotd})$ (2)	Firm's within origin mkt share $\ln(ms_{fiotd})$ (3)	Origin's mkt share $\ln(ms_{ioidt})$ (4)	No. of firms (PPML) (5)
Intermediate Goods					
Tariff _{ioidt}	0.02 (0.369)	0.10 (0.107)	1.76*** (0.628)	-1.51*** (0.446)	-1.48*** (0.229)
PTA _{oidt}	-0.03 (0.028)	-0.03* (0.014)	-0.16*** (0.037)	0.03 (0.034)	0.03** (0.011)
R ²	0.69	0.90	0.77	0.75	-
Observations	5,777,419	5,735,168	5,777,419	522,910	1,384,137
Highly Differentiated Intermediate Goods					
Tariff _{ioidt}	-0.50 (1.292)	0.21 (0.389)	-1.98 (1.342)	-1.20 (1.254)	-0.77 (0.499)
PTA _{oidt}	0.08 (0.079)	-0.01 (0.050)	-0.10 (0.088)	0.14 (0.089)	0.03 (0.025)
R ²	0.64	0.91	0.75	0.80	-
Observations	693,697	689,152	693,697	57,169	143,809
Less Differentiated Intermediate Goods					
Tariff _{ioidt}	-0.19 (0.366)	0.11 (0.108)	2.00*** (0.684)	-1.67*** (0.477)	-1.52*** (0.252)
PTA _{oidt}	-0.02 (0.030)	-0.03*** (0.015)	-0.13*** (0.038)	0.01 (0.037)	0.02 (0.013)
R ²	0.71	0.90	0.77	0.74	-
Observations	4,692,853	4,656,161	4,692,853	448,496	1,201,198
Fixed Effects					
Firm-product-origin-year	✓	✓	✓		
Product-origin-year				✓	✓
Product-destination-year	✓	✓	✓	✓	✓
Origin-destination	✓	✓	✓	✓	✓

Notes: The dependent variable is the firm's log (tariff-inclusive) export value in column (1), the firm's log (tariff-exclusive) unit value in column (2), the log of the firm's share of its country's trade with the destination in column (3), the log of the country's (tariff-inclusive) export value to the destination market in column (4) and the number of firms in column (5). Products are separated into different groups based on the CCHS and BEC (Revision 4) classification systems. Standard errors, reported in parentheses, are clustered at the product-destination level, and we denote statistical significance with *** p<0.01, ** p<0.05, and * p<0.1. Estimates are based on an integrated dataset of firms' exports from eleven countries built from the World Bank Exporter Dynamics Database, China's Customs Authority, and Egypt's Customs Authority, as well as tariff data from the WTO and Feenstra and Romalis (2014), and the World Bank Deep Trade Agreements Database.

participation is also particularly pronounced for this set of products. Finally, the bottom two panels show that tariff liberalizations have little effect on less differentiated consumption goods.

The overall pattern is consistent with the idea that the markets for consumption goods are oligopolistic, and have a larger within-origin elasticity of substitution σ relative to their across-origin elasticity of substitution ρ . As a result, firms which sell consumption goods, and particularly highly differentiated consumption goods, seem to respond strongly to additional entry due to trade liberalizations by lowering their markups.

5.3 Trade policy and destination income levels

The previous two subsections have employed partitions of product space. In this subsection, we take an alternative approach and split our sample based on income-levels in the destination to isolate markets with different scopes for market power. Rich countries have more power to enforce idiosyncratic national policies such as health or safety regulations that create an international segmentation of markets. They also have larger and more diverse markets, and a larger proportion of their imports are highly differentiated.³⁰ We would therefore expect trade policy to have larger effects on markups in high-income countries. To investigate this idea, we split the destinations in our sample into three groups, high-, middle- and low-income countries, based on their per-capita income levels in 1999 and repeat our analysis for each of these three subsamples.

Table 6 presents our estimates. The results show that the tariff elasticities of firms' market shares and markups are largest in high income destinations. A 10% tariff reduction leads to a 16% increase in firms' market shares in high income destinations, but has little effect on market shares in middle or low income destinations. At the same time, it decreases markups by 8% in high income destinations, by 2% in middle income destinations, and has no effect on markups in low income destinations. Our findings may reflect that countries with large markets where firms from many different countries compete can support consumption of many different varieties of a product, and that oligopolistic competition between firms which produce varieties distinct to their origin is more plausible. As a result, the within-origin elasticity of substitution σ is larger relative to the cross-origin elasticity of substitution ρ and the within-origin reallocation effect is more likely to dominate in destinations with higher incomes.

³⁰See appendix A9 for details.

Table 6: Trade policy elasticities by destination country income level

	Firm's mkt share in the dest. $\ln(\omega_{fiodt})$ (1)	Markups $\ln(\mu_{fiodt})$ (2)	Firm's within origin mkt share $\ln(ms_{fiodt})$ (3)	Origin's mkt share $\ln(ms_{iodt})$ (4)	No. of firms (PPML) (5)
High Income Countries					
Tariff _{iodt}	-1.60** (0.637)	0.83*** (0.143)	2.35*** (0.594)	-1.21* (0.709)	-1.87*** (0.231)
PTA _{odt}	0.13** (0.052)	-0.06*** (0.016)	0.06 (0.057)	-0.07 (0.044)	0.15*** (0.015)
R ²	0.65	0.90	0.78	0.82	-
Observations	8,257,917	8,231,061	8,257,917	487,476	1,138,367
Middle Income Countries					
Tariff _{iodt}	-0.24 (0.203)	0.21*** (0.077)	2.94*** (0.418)	-2.70*** (0.287)	-2.04*** (0.162)
PTA _{odt}	-0.02 (0.023)	-0.01 (0.012)	-0.04 (0.036)	-0.02 (0.030)	-0.02** (0.010)
R ²	0.74	0.91	0.82	0.75	-
Observations	4,300,329	4,277,859	4,300,329	462,975	1,207,037
Low Income Countries					
Tariff _{iodt}	0.46 (0.845)	0.03 (0.427)	0.85 (0.886)	-0.80 (0.870)	-1.36*** (0.277)
PTA _{odt}	-0.43*** (0.165)	0.01 (0.084)	0.12 (0.165)	-0.21 (0.157)	-0.38*** (0.041)
R ²	0.76	0.93	0.85	0.81	-
Observations	867,236	865,706	867,236	46,506	180,981
Fixed Effects					
Firm-product-origin-year	✓	✓	✓		
Product-origin-year				✓	✓
Product-destination-year	✓	✓	✓	✓	✓
Origin-destination	✓	✓	✓	✓	✓

Notes: The dependent variable is the firm's log (tariff-inclusive) export value in column (1), the firm's log (tariff-exclusive) unit value in column (2), the log of the firm's share of its country's trade with the destination in column (3), the log of the country's (tariff-inclusive) export value to the destination market in column (4) and the number of firms in column (5). Countries are separated into high income, middle income and low income destinations according to World Bank lending groups in 1999. Standard errors, reported in parentheses, are clustered at the product-destination level, and we denote statistical significance with *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$. Estimates are based on an integrated dataset of firms' exports from eleven countries built from the World Bank Exporter Dynamics Database, China's Customs Authority, and Egypt's Customs Authority, as well as tariff data from the WTO and [Feenstra and Romalis \(2014\)](#), and the World Bank Deep Trade Agreements Database.

6 Conclusion

Understanding the welfare implications of trade agreements has long been a central focus of the international economics literature. The role of competition and markup adjustments, and the question of whether the pro-competitive gains from trade are elusive, is at the core of recent debates.³¹ Despite several theoretical contributions, empirical evidence on how foreign exporters compete and adjust their markups to trade policy changes in a multi-country world remains scarce.

In this paper, we exploit product exports by firms from eleven emerging economies and investigate how tariffs and preferential trade agreements affect the ways in which firms compete and the markups they charge. We find, surprisingly, that in response to a bilateral tariff cut, foreign exporters lower their markups while their market share in the destination increases – an observation that contradicts the predictions of standard oligopolistic competition models.

We show this puzzling empirical finding can be rationalized theoretically in a more general multi-country framework that allows for a different degree of oligopolistic competition within and across origins. Our theoretical model suggests *two* market share reallocation effects matter for firms’ markup adjustments after a bilateral tariff cut: (1) a “within-origin” reallocation effect that reduces the firms’ optimal markups and (2) an “across-origin” reallocation effect that increases the firms’ optimal markups. We show that firms’ within-origin market share changes have larger impacts on their markups than the across-origin market share changes when the varieties firms produce are more substitutable within an origin than across origins. Empirically, we find strong support for these two effects.

A key feature of our theoretical framework is that we do not need to take a stand on the magnitude of the two reallocation effects *ex ante*. Both reallocation effects exist in our triple-nested demand framework and the overall markup adjustments and their associated pro-competitive effects depend on which of the two effects dominates. While our empirical results suggest the “within-origin reallocation effect” dominates for the eleven low and middle income countries in our study over the period we examine, it is possible that these two effects cancel each other out and that there is no change in the average markups of exporters in

³¹See [Bagwell and Staiger \(2016\)](#) for a comprehensive review of the theoretical literature on the welfare consequences of trade agreements and [Ossa \(2016\)](#) for a summary of the literature on quantitative modelling of trade agreements. While early contributions investigated the efficiency properties of trade agreements under perfect competition ([Bagwell and Staiger \(1999\)](#)), more recent studies have examined welfare impacts under more complex market structures featuring price formation under bilateral bargaining ([Antràs and Staiger \(2012\)](#)) or in an environment with variable markups ([Bagwell and Lee \(2020\)](#)). See [Edmond, Midrigan and Xu \(2015\)](#) and [Arkolakis, Costinot, Donaldson and Rodríguez-Clare \(2018\)](#) for recent debates on the pro-competitive gains of trade.

many other cases. In this vein, our framework provides a possible theoretical explanation for the lack of export price adjustments by Chinese exporters during the US-China Trade War (see e.g., [Amiti, Redding and Weinstein \(2019\)](#), [Fajgelbaum, Goldberg, Kennedy and Khandelwal \(2020\)](#), [Cavallo, Gopinath, Neiman and Tang \(2021\)](#)). The lack of any price response to tariff hikes could arise if the two reallocation effects discussed in our paper have offsetting impacts on prices.

References

- Amiti, Mary, and Jozef Konings.** 2007. “Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia.” *American Economic Review*, 97(5): 1611–1638.
- Amiti, Mary, Oleg Itskhoki, and Jozef Konings.** 2019. “International Shocks, Variable Markups, and Domestic Prices.” *The Review of Economic Studies*, 86(6): 2356–2402.
- Amiti, Mary, Stephen J. Redding, and David E. Weinstein.** 2019. “The Impact of the 2018 Tariffs on Prices and Welfare.” *Journal of Economic Perspectives*, 33(4): 187–210.
- Anderson, James E., and Eric van Wincoop.** 2003. “Gravity with Gravitas: A Solution to the Border Puzzle.” *American Economic Review*, 93(1): 170–192.
- Antràs, Pol, and Robert W. Staiger.** 2012. “Offshoring and the Role of Trade Agreements.” *American Economic Review*, 102(7): 3140–83.
- Arkolakis, Costas, Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare.** 2018. “The Elusive Pro-Competitive Effects of Trade.” *The Review of Economic Studies*, 86(1): 46–80.
- Atkeson, Andrew, and Ariel Burstein.** 2008. “Pricing-to-Market, Trade Costs, and International Relative Prices.” *The American Economic Review*, 98(5): 1998–2031.
- Bagwell, K., and R.W. Staiger.** 2016. “Chapter 8 - The Design of Trade Agreements.” In . Vol. 1 of *Handbook of Commercial Policy*, ed. Kyle Bagwell and Robert W. Staiger, 435–529. North-Holland.
- Bagwell, Kyle, and Robert W. Staiger.** 1999. “An Economic Theory of GATT.” *The American Economic Review*, 89(1): 215–248.
- Bagwell, Kyle, and Seung Hoon Lee.** 2020. “Trade policy under monopolistic competition with firm selection.” *Journal of International Economics*, 127: 103379.
- Baier, Scott L., and Jeffrey H. Bergstrand.** 2007. “Do free trade agreements actually increase members’ international trade?” *Journal of International Economics*, 71(1): 72 – 95.
- Baier, Scott L., Jeffrey H. Bergstrand, and Michael Feng.** 2014. “Economic integration agreements and the margins of international trade.” *Journal of International Economics*, 93(2): 339 – 350.

- Bas, Maria, Thierry Mayer, and Mathias Thoenig.** 2017. “From micro to macro: Demand, supply, and heterogeneity in the trade elasticity.” *Journal of International Economics*, 108: 1–19.
- Bastos, Paulo, and Joana Silva.** 2010. “The quality of a firm’s exports: Where you export to matters.” *Journal of International Economics*, 82(2): 99 – 111.
- Bernard, Andrew B., Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum.** 2003. “Plants and Productivity in International Trade.” *American Economic Review*, 93(4): 1268–1290.
- Bortoluzzi, Aldo Pazzini, Ana Margarida Fernandes, and Martha Denisse Pierola.** 2015. “Exporter Dynamics Database – Version 2.0: A Technical Note.” *World Bank Technical Note*.
- Bown, Chad P., and Meredith A. Crowley.** 2006. “Policy Externalities: How US Antidumping Affects Japanese Exports to the EU.” *European Journal of Political Economy*, 22(1): 696–714.
- Brander, James, and Paul Krugman.** 1983. “A ‘reciprocal dumping’ model of international trade.” *Journal of International Economics*, 15(3): 313–321.
- Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang.** 2021. “Tariff Pass-Through at the Border and at the Store: Evidence from US Trade Policy.” *American Economic Review: Insights*, 3(1): 19–34.
- Cebeci, Tolga.** 2015. “A Consistent Concordance Among Harmonized System 1996, 2002, 2007, and 2012 Classifications.” *World Bank Technical Note*.
- Cebeci, Tolga, Ana Margarida Fernandes, Caroline Freund, and Martha Denisse Pierola.** 2012. “Exporter Dynamics Database.” *Policy Research Working Paper 6229*.
- Corsetti, Giancarlo, Meredith Crowley, and Lu Han.** 2022. “Invoicing and the dynamics of pricing-to-market: Evidence from UK export prices around the Brexit referendum.” *Journal of International Economics*, 135: 103570.
- Corsetti, Giancarlo, Meredith Crowley, Lu Han, and Huasheng Song.** 2021. “Markets and Markups : A New Empirical Framework and Evidence on Exporters from China.” *CEPR Discussion Paper*, 13904.
- De Blas, Beatriz, and Katheryn N. Russ.** 2015. “Understanding Markups in the Open Economy.” *American Economic Journal: Macroeconomics*, 7(2): 157–180.

- De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik.** 2016. “Prices, Markups, and Trade Reform.” *Econometrica*, 84(2): 445–510.
- Eaton, Jonathan, and Gene M. Grossman.** 1986. “Optimal Trade and Industrial Policy under Oligopoly.” *The Quarterly Journal of Economics*, 101(2): 383–406.
- Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu.** 2015. “Competition, markups, and the gains from international trade.” *American Economic Review*, 105(10): 3183–3221.
- Fajgelbaum, Pablo D, Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal.** 2020. “The Return to Protectionism.” *The Quarterly Journal of Economics*, 135(1): 1–55.
- Feenstra, Robert C.** 2004. *Advanced international trade: theory and evidence*. Princeton, NJ:Princeton University Press.
- Feenstra, Robert C., and John Romalis.** 2014. “International Prices and Endogenous Quality.” *The Quarterly Journal of Economics*, 129(2): 477–527.
- Fitzgerald, Doireann, and Stefanie Haller.** 2014. “Pricing-to-Market: Evidence from Plant-Level Prices.” *The Review of Economic Studies*, 81(2): 761–786.
- Fitzgerald, Doireann, and Stefanie Haller.** 2018. “Exporters and Shocks.” *Journal of International Economics*, 113: 154–171.
- Head, Keith, and Thierry Mayer.** 2014. “Chapter 3 - Gravity Equations: Workhorse, Toolkit, and Cookbook.” In *Handbook of International Economics*. Vol. 4 of *Handbook of International Economics*, ed. Gita Gopinath, Elhanan Helpman and Kenneth Rogoff, 131 – 195. Elsevier.
- Helpman, Elhanan, and Paul Krugman.** 1985. *Market Structure and Foreign Trade*. Cambridge, MA:MIT Press.
- Hummels, David, and Alexandre Skiba.** 2004. “Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture.” *Journal of Political Economy*, 112(6): 1384–1402.
- Kikkawa, Ayuma Ken, Yuan Mei, and Pablo Robles Santamarina.** 2019. “The Impact of NAFTA on Prices and Competition: Evidence from Mexican Manufacturing Plants.”

- Konings, Jozef, and Hylke Vandenbussche.** 2005. “Antidumping protection and markups of domestic firms.” *Journal of International Economics*, 65(1): 151 – 165.
- Markusen, James R., and Anthony J. Venables.** 1988. “Trade policy with increasing returns and imperfect competition: Contradictory results from competing assumptions.” *Journal of International Economics*, 24(3): 299–316.
- Melitz, Marc J.** 2003. “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity.” *Econometrica*, 71(6): 1695–1725.
- Melitz, Marc J., and Gianmarco I. P. Ottaviano.** 2008. “Market Size, Trade, and Productivity.” *The Review of Economic Studies*, 75(1): 295–316.
- Ossa, R.** 2016. “Chapter 4 - Quantitative Models of Commercial Policy.” In . Vol. 1 of *Handbook of Commercial Policy*, ed. Kyle Bagwell and Robert W. Staiger, 207–259. North-Holland.
- Pierce, Justin R.** 2011. “Plant-level responses to antidumping duties: Evidence from U.S. manufacturers.” *Journal of International Economics*, 85(2): 222 – 233.
- Rauch, James E.** 1999. “Networks Versus Markets in International Trade.” *Journal of International Economics*, 48(1): 7–35.
- Redding, Stephen, and Anthony J. Venables.** 2004. “Economic geography and international inequality.” *Journal of International Economics*, 62(1): 53 – 82.
- Viner, Jacob.** 1950. *The Customs Union Issue*. New York: Carnegie Endowment for International Peace.

A Data appendix

A.1 Feenstra and Romalis (2014)

We augment our tariff data with information on the phase-in of trade agreements contained in the dataset created by [Feenstra and Romalis \(2014\)](#), who compile data on MFN and preferential tariffs between 1984 and 2011 for a large number of countries at the 4-digit SITC (Revision 2) product level. They impute preferential tariffs by extracting information on the phase-in of preferential tariffs from more than 100 trade agreements, which they express as fractions of the MFN tariff and then multiply with the MFN tariffs in their dataset.

To recover data on the phase-in of trade agreements from this dataset, which is available from Robert Feenstra’s website, we express preferential tariffs as a fraction of MFN tariffs, calculate the mode across products within a country-pair and year and then consider the fraction of products with a value equal to that mode. We drop country-pair-years with more than one mode, country-pair-years where fewer than 25% of industries have values equal to the mode and country-pair-years where more than 25% of industries have missing values. For the remaining country-pair-years, we imputed missing preferential tariffs in our dataset by multiplying our MFN tariff with the fraction we get from this procedure.

A.2 Consolidated product codes

We consolidate HS codes to ensure that the product codes in our analysis are consistent over time. Our trade, tariff and commodity classification data are reported based on the HS product classification system. Since our data span a large number of years and the HS system is updated periodically, our data could feature up to four different revisions of the HS system (HS1996, HS2002, HS2007 and HS2012). We transform HS codes into consolidated HS codes which are constant over time by identifying networks of related product codes and assigning a unique consolidated code to each network, similar to [Cebeci \(2015\)](#). This reduces the number of distinct products in the HS system from 6,293 to 4,039.

A.3 Product differentiation

We determine the degree of product differentiation for different products in our dataset by using the CCHS commodity classification system, which sorts products into highly and less differentiated goods. This classification is based on the fact that there are a large number of *measure words* in the Chinese language, and that the choice of measure word used for a given product is predetermined by Chinese grammar and linguistics and therefore reflects a good’s intrinsic physical features. The core idea here is that goods whose quantity is

recorded in specific countable units, such as motorcycles or consumer electronics, are more differentiated than goods whose quantity is recorded in continuous units, such as canned tomato paste or industrial chemicals. In Chinese trade data, quantity is reported in more than 30 indigenous Chinese units of measure, including distinct words representing the unit count of products such as wheeled vehicles, engines, and upper-body clothing articles, as well as more general terms for weight or volume. The CCHS classification exploits this distinction between what linguists refer to as count and mass measure words to construct a general product classification for the Harmonized System.³²

The CCHS classification is available for more than 4,800 products at the HS6 level. In the 10 instances in which the CCHS code of a product is not the same for all HS6 codes within a consolidated HS code, we set the CCHS code to missing and disregard such products in our analysis.

A.4 Broad Economic Categories

To further refine the product space, we use information on product end-use categories provided by the UN's Broad Economic Categories classification (Revision 4) to distinguish between intermediate and consumption goods. This information is available via the UN Statistics Division's correspondence tables, which map HS6 codes to BEC categories and allow us to create a mapping from consolidated HS codes to BEC categories. There are only 17 consolidated codes with HS6 codes that correspond to more than one BEC category, and we omit these cases from our analysis.

³²See Corsetti, Crowley, Han and Song (2021) for a more extensive discussion of measure words and evidence of how they are used in other East Asian customs records.

Table A7: Summary statistics of our estimation sample

	Mean	S.D.	Percentile					Observations
			1 st	25 th	Median	75 th	99 th	
Trade value, $p_{fi}y_{fi}$	304,274	19,635,593	14	2,290	11,304	50,885	3,102,922	15,853,618
Tariff exclusive FOB price, p_{fi}^b	982	312,274	.046	1	2.6	8.2	3,999	15,793,386
Firm's within-origin market share, ms_{fi}	0.073	0.19	.00000055	0.00044	0.004	0.32	1	15,853,618
Origin's trade value in a destination-product market, $\sum_{f \in \mathcal{F}_{i}^{odt}} p_{fi}y_{fi}$	4,881,001	206,237,116	3.7	1,796	23,691	245,585	54,653,516	1,067,240
Number of firms, N_{i}^{odt}	4	36	0	0	0	1	65	6,547,520
Bilateral Ad Valorem Tariff (percent), $(\tau_{i}^{odt} - 1) * 100$	7.8	9.9	0	1.6	5	12	40	15,853,618
Preferential Trade Agreement Dummy, PTA_{odt}	.11	.31	0	0	0	0	1	15,853,618

Notes: This table presents summary statistics for the variables used in the estimation specifications reported in Table 2. The estimation sample is smaller than the full sample of 25,176,098 observations because fixed effects estimation excludes “singleton” observations. The row for “Number of firms” presents statistics for the sample including zero trade flows. We create a zero trade flow for year t whenever we see at least one firm exporting a product i from an origin o to a destination d in any year for which we observe o 's exports but not in year t .

Table A8: The number of exporters

	Mean	25th Percentile	Median	75th Percentile	Observations
Initial Sample					
Number of Firms	8.89	1.00	2.00	5.00	2,956,796
Number of Entrants	6.27	1.00	1.00	4.00	2,403,979
Number of Exiters	6.12	1.00	2.00	4.00	1,744,997
Number of Incumbents	1.87	0.00	0.00	1.00	2,000,356
Full Sample					
Number of Firms	4.01	0.00	0.00	1.00	6,547,520
Number of Entrants	2.85	0.00	0.00	1.00	5,284,851
Number of Exiters	4.81	1.00	1.00	3.00	2,383,130
Number of Incumbents	0.78	0.00	0.00	0.00	4,820,554

Notes: This table presents summary statistics for i) the number of firms from an origin o selling product i to destination d at time t , ii) the number of firms that did not sell this product to that destination in period $t - 1$ but do so in period t , iii) the number of firms that sell this product to that destination in period t but do not do so in period $t + 1$ and iv) the number of firms that sell this product to that destination in periods $t - 1$, t and $t + 1$. We create a zero trade flow for year t whenever we see a firm f exporting a product i from an origin o to a destination d in any year for which we observe o 's exports but not in year t . The first panel presents statistics calculated for the sample excluding zero trade flows, the second for the sample including zero trade flows.

Table A9: Decomposition of products sold to countries of different income levels

Highly Differentiated Goods		Less Differentiated Goods		Observations
Consumption	Intermediate	Consumption	Intermediate	
<i>High Income Countries</i>				
27.89	4.01	16.84	30.30	14,982,886
<i>Middle Income Countries</i>				
22.97	4.69	15.42	34.49	8,651,823
<i>Low Income Countries</i>				
12.99	4.99	13.03	44.63	2,607,394

Notes: This table presents the proportion of observations that fall into the four different product categories, broken down by destination income.

B Supplementary model results

B.1 Deriving the demand elasticity under our triple-nested CES demand framework

Upon entry, the operational profit of the firm is given by

$$\begin{aligned}\pi_{fiodt}^{operational} &= \left(\frac{p_{fiodt}}{\tau_{fiodt}} - mc_{fiodt} \right) y_{fiodt} \\ &= \left(\frac{p_{fiodt}}{\tau_{fiodt}} - mc_{fiodt} \right) \alpha_{fiodt} p_{fiodt}^{-\sigma} (p_{fiodt})^{\sigma-\rho} (p_{idt})^{\rho-\eta} P_{dt}^{\eta} Y_{dt}\end{aligned}$$

Maximizing profits with respect to p_{fiodt} , yields the first order condition which can be rearranged to get:

$$\frac{y_{fiodt}}{\tau_{fiodt}} + \left(\frac{p_{fiodt}}{\tau_{fiodt}} - mc_{fiodt} \right) \frac{\partial y_{fiodt}}{\partial p_{fiodt}} = 0$$

Define the price elasticity of demand as

$$\varepsilon_{fiodt} \equiv - \frac{\partial y_{fiodt} p_{fiodt}}{\partial p_{fiodt} y_{fiodt}}$$

For a given ε_{fiodt} , the optimal price can be easily derived and is given by (5). The tricky part, however, is to calculate the demand elasticity ε_{fiodt} , which can be expressed as

$$\varepsilon_{fiodt} = - \frac{\partial [p_{fiodt}^{-\sigma} (p_{fiodt})^{\sigma-\rho} (p_{idt})^{\rho-\eta}]}{\partial p_{fiodt}} \frac{p_{fiodt}}{p_{fiodt}^{-\sigma} (p_{fiodt})^{\sigma-\rho} (p_{idt})^{\rho-\eta}}$$

We now calculate the elements of the demand elasticity one-by-one using the chain rule:

$$\frac{\partial [p_{fiodt}^{-\sigma} (p_{fiodt})^{\sigma-\rho} (p_{idt})^{\rho-\eta}]}{\partial p_{fiodt}} = \frac{\partial [p_{fiodt}^{-\sigma} (p_{fiodt})^{\sigma-\rho}]}{\partial p_{fiodt}} (p_{idt})^{\rho-\eta} + p_{fiodt}^{-\sigma} (p_{fiodt})^{\sigma-\rho} \frac{\partial [(p_{idt})^{\rho-\eta}]}{\partial p_{fiodt}}$$

$$\frac{\partial [p_{fiodt}^{-\sigma} (p_{fiodt})^{\sigma-\rho}]}{\partial p_{fiodt}} = -\sigma (p_{fiodt})^{-\sigma-1} (p_{fiodt})^{\sigma-\rho} + p_{fiodt}^{-\sigma} \frac{\partial [(p_{fiodt})^{\sigma-\rho}]}{\partial p_{fiodt}}$$

$$\begin{aligned}
\frac{\partial [(p_{iodt})^{\sigma-\rho}]}{\partial p_{fiodt}} &= \frac{\partial \left(\sum_{f \in \mathcal{F}_{iodt}} \alpha_{fiodt} p_{fiodt}^{1-\sigma} \right)^{\frac{\sigma-\rho}{1-\sigma}}}{\partial p_{fiodt}} \\
&= (\sigma - \rho) \alpha_{fiodt} p_{fiodt}^{-\sigma} \left(\sum_{f \in \mathcal{F}_{iodt}} \alpha_{fiodt} p_{fiodt}^{1-\sigma} \right)^{\frac{\sigma-\rho}{1-\sigma}-1} \\
&= (\sigma - \rho) \alpha_{fiodt} p_{fiodt}^{-\sigma} (p_{iodt})^{2\sigma-\rho-1}
\end{aligned}$$

$$\begin{aligned}
\frac{\partial [(p_{id})^{\rho-\eta}]}{\partial p_{fiodt}} &= \frac{\partial \left[\left(\sum_{o \in \mathcal{H}} p_{iodt}^{1-\rho} \right)^{\frac{\rho-\eta}{1-\rho}} \right]}{\partial p_{fiodt}} \\
&= \frac{\partial \left[\left(\sum_{o \in \mathcal{H}} p_{iodt}^{1-\rho} \right)^{\frac{\rho-\eta}{1-\rho}} \right]}{\partial p_{iodt}} \frac{\partial p_{iodt}}{\partial p_{fiodt}} \\
&= (\rho - \eta) p_{iodt}^{-\rho} (p_{id})^{2\rho-\eta-1} \cdot \alpha_{fiodt} p_{fiodt}^{-\sigma} (p_{iodt})^{\sigma} \\
&= \alpha_{fiodt} (\rho - \eta) p_{fiodt}^{-\sigma} (p_{iodt})^{\sigma-\rho} (p_{id})^{2\rho-\eta-1}
\end{aligned}$$

$$\begin{aligned}
\frac{\partial [p_{fiodt}^{-\sigma} (p_{iodt})^{\sigma-\rho} (p_{id})^{\rho-\eta}]}{\partial p_{fiodt}} &= \frac{\partial [p_{fiodt}^{-\sigma} (p_{iodt})^{\sigma-\rho}]}{\partial p_{fiodt}} (p_{id})^{\rho-\eta} + p_{fiodt}^{-\sigma} (p_{iodt})^{\sigma-\rho} \frac{\partial [(p_{id})^{\rho-\eta}]}{\partial p_{fiodt}} \\
&= -\sigma (p_{fiodt})^{-\sigma-1} (p_{iodt})^{\sigma-\rho} (p_{id})^{\rho-\eta} + (\sigma - \rho) \alpha_{fiodt} p_{fiodt}^{-2\sigma} (p_{iodt})^{2\sigma-\rho-1} (p_{id})^{\rho-\eta} \\
&\quad + \alpha_{fiodt} (\rho - \eta) p_{fiodt}^{-2\sigma} (p_{iodt})^{2\sigma-2\rho} (p_{id})^{2\rho-\eta-1} \\
&= -p_{fiodt}^{-\sigma-1} (p_{iodt})^{\sigma-\rho} (p_{id})^{\rho-\eta} [\sigma - (\sigma - \rho) m s_{fiodt} - (\rho - \eta) m s_{fiodt} m s_{iodt}]
\end{aligned}$$

Using the above relationships, we can express the demand elasticity as a function of market shares:

$$\begin{aligned}
\varepsilon_{fiodt} &= \sigma - (\sigma - \rho) m s_{fiodt} - (\rho - \eta) m s_{fiodt} m s_{iodt} \\
&= \sigma - m s_{fiodt} [\sigma - \rho + (\rho - \eta) m s_{iodt}]
\end{aligned}$$

B.2 Proof of Proposition 1

Proof. Using equations (5) and (6), it can be shown that:

$$\widehat{\mu}_{fiodt} = -\frac{1}{\varepsilon_{fiodt} - 1} \widehat{\varepsilon}_{fiodt} \quad (18)$$

$$\widehat{\varepsilon}_{fiodt} = -\frac{\sigma - \varepsilon_{fiodt}}{\varepsilon_{fiodt}} \widehat{m}s_{fiodt} - \frac{(\rho - \eta)m_s_{fiodt}m_s_{iodt}}{\varepsilon_{fiodt}} \widehat{m}s_{iodt} \quad (19)$$

Substituting (19) into (18), we get

$$A(\cdot) \equiv \frac{\sigma - \varepsilon_{fiodt}}{\varepsilon_{fiodt}(\varepsilon_{fiodt} - 1)}; \quad B(\cdot) \equiv \frac{(\rho - \eta)m_s_{fiodt}m_s_{iodt}}{\varepsilon_{fiodt}(\varepsilon_{fiodt} - 1)} \quad (20)$$

Since $\varepsilon_{fiodt}(\varepsilon_{fiodt} - 1)$ is strictly larger than 0, the sign of $A(\cdot) - B(\cdot)$ depends on the sign of

$$\sigma - \varepsilon_{fiodt} - (\rho - \eta)m_s_{fiodt}m_s_{iodt} = (\sigma - \rho)m_s_{fiodt} \quad (21)$$

Given that $m_s_{fiodt} > 0$, $A(\cdot) - B(\cdot) > 0$ iff $\sigma > \rho$ and $A(\cdot) - B(\cdot) = 0$ iff $\sigma = \rho$. \square

B.3 Further discussion of figure 1 patterns

To understand the patterns of the top two panels of figure 1, it is useful to consider two additional micro channels. First, there is a general *substitutability* channel. As goods become more substitutable within and across origins, firms have less market power and make smaller markup adjustments. This substitutability channel is most transparent in the diagonal elements of the two panels. When $\rho = \sigma$, our model collapses to the [Atkeson and Burstein \(2008\)](#) model and the markup adjustment decreases as the substitutability across products (both within and across origins) increases.

Second, there is a *relevance* channel when we move the within-origin elasticity of substitution σ away from the across-origin elasticity of substitution ρ – different values of σ and ρ change the relative importance of a specific market share change. When σ increases relative to ρ , changes in the firm's within-origin market share $\widehat{m}s_{fiodt}$ become more important and its markup responds more to changes in this market share. In contrast, when ρ increases relative to σ , changes in the origin's market share in the destination $\widehat{m}s_{iodt}$ become more important and the firm's markup responds more to the changes in the across-origin market share.

With these two channels in mind, we are ready to analyse the patterns of the off-diagonal elements of figure 1. There are two patterns which stand out in the top-left panel. First, as we increase the within-origin elasticity of substitution σ for a given value of the across-origin

elasticity of substitution ρ , the firm's markup adjustments initially increase, but eventually decline. This is because the substitutability and relevance channels work in opposite directions. We can see the relevance effect initially dominates for lower levels of σ , and the substitutability effect only prevails once σ reaches a certain level. Second, as we increase the across-origin elasticity of substitution ρ for a given level of the within-origin-elasticity of substitution σ , the firm's markup adjustment in response to a 1% change in the within-origin market share is decreasing. This is because both channels now work in the same direction. The firm's goods have become more substitutable, and changes in the firm's within-origin market share are less relevant.

Similarly, for the off-diagonal elements of the top right panel, we can identify two broad patterns. First, as we increase ρ for a given value of σ , the firm's markup response is increasing. This is because the relevance channel dominates the substitutability channel for the parameter values shown in the figure. We can see, however, that the rate at which the firm's markup response is increasing as we increase ρ is declining as ρ increases and the substitutability effect becomes more important. Second, as we increase σ for a given value of ρ , the firm's markup response is decreasing. This is because both channels are now working in the same direction and reduce the magnitude of the markup adjustment: (i) the substitutability channel suggests a higher σ implies lower market power and smaller markup adjustments and (ii) the relevance channel suggests a higher σ means that changes in the firm's within-origin market share (the origin's market share in the destination) are relatively more (less) relevant.

B.4 Within- and across-origin reallocation effects for firms and origins with different initial market shares

Figure B2 shows values of the $A(\cdot)$ and $B(\cdot)$ functions fixing the three elasticities (i.e., σ, ρ, η) while varying the initial market shares of firms and origins. The top-left panel of B2 shows the markup adjustments for a 1% increase in the firm's within-origin market share while keeping the origin's market share in the destination fixed ($\widehat{m}s_{fiodt} = 1\%$ and $\widehat{m}s_{iodt} = 0\%$). In contrast, the top-right panel of B2 shows the markup adjustments for a 1% increase in the origin's market share in the destination while keeping the firm's within-origin market share fixed ($\widehat{m}s_{fiodt} = 0\%$ and $\widehat{m}s_{iodt} = 1\%$). The x-axis in each figure measures the firm's initial within-origin market share $m s_{fiodt}$ and the y-axis measures the origin's initial market share in the destination $m s_{iodt}$.

We can see from the top two figures that markup adjustments are larger as the two initial market shares ($m s_{fiodt}$ and $m s_{iodt}$) increase. This is a very intuitive result. As a firm or a

origin becomes more important, a 1% change in its market share will have a much bigger impact on the market structure and thus the firm (and its competitors) make larger markup adjustments.

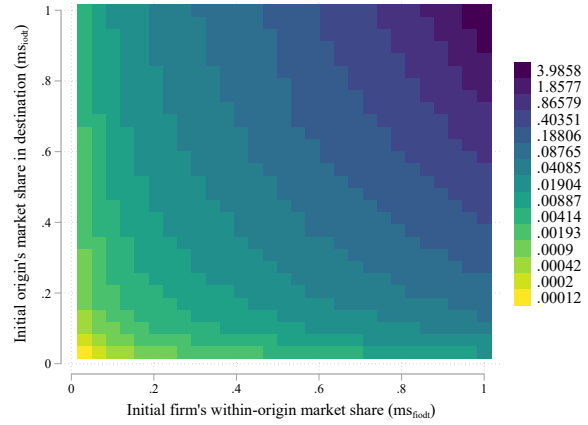
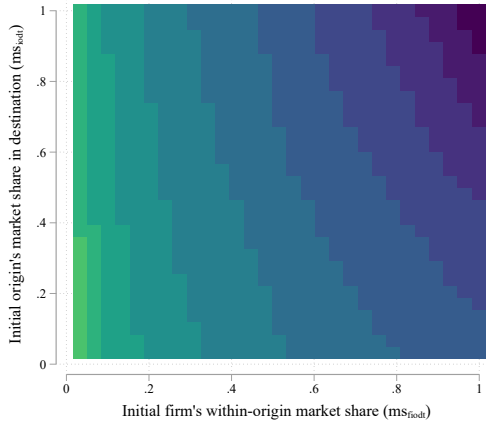
The bottom left panel shows markup adjustments for a 1% increase in the firm's within-origin market share and a 1% drop in the origin's market share in the destination, which together keep the firm's overall market share in the destination constant ($\widehat{m}s_{fiott} = 1\%$, $\widehat{m}s_{iodt} = -1\%$ and $\widehat{\omega}_{fiott} = \widehat{m}s_{fiott} + \widehat{m}s_{iodt} = 0\%$). As predicted by Proposition 1, we find the within origin reallocation effect dominates when $\sigma > \rho$.

The bottom-right panel shows the ratio of the two functions (i.e., $A(\cdot)/B(\cdot)$). An interesting feature we find in this panel is that the ratio tends to be significantly larger for origins with relatively small market shares in the destination. This implies that the within-origin reallocation effect is more likely to dominate the direction of markup adjustments for origins with small market shares. It is worth noting that this happens to be the case for our data since the origin's market share in the destination is small for most industries for the eleven origin countries in our dataset (except perhaps for China). While it is possible that future studies, which investigate different datasets, will find different (and maybe weaker) markup adjustments, the two opposing reallocation effects we consider should remain important and could explain exporters markup responses even in cases where standard oligopolistic competition models cannot.

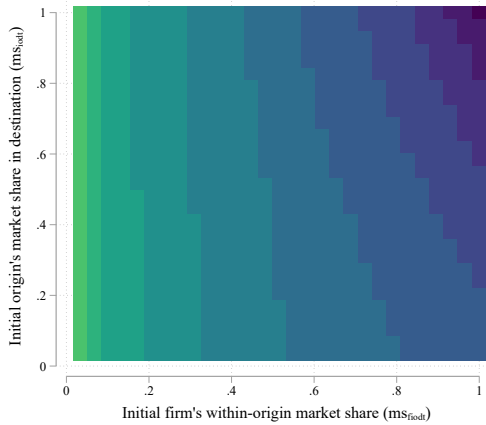
Figure B2: Visualizing the within- and across-origin reallocation effects on markups varying market shares while fixing $\sigma = 4.0$, $\rho = 2.5$, and $\eta = 1.2$

(A) Within-origin reallocation effect (for a 1% change in the firm's within-origin market share i.e., $\widehat{ms}_{fiobt} = 1\%$)

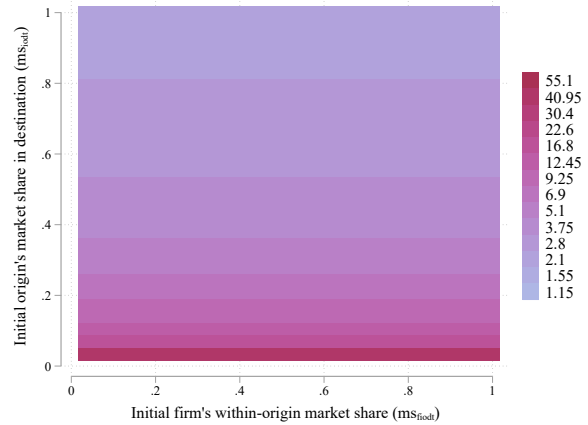
(B) Across-origin reallocation effect (for a 1% change in the origin's market share in the destination, i.e., $\widehat{ms}_{iobt} = 1\%$)



(A - B) Level difference of the two effects



(A/B) Ratio of the two effects



Notes: The above figures show the values of $A(\sigma, \rho, \eta, ms_{fiobt}, ms_{iobt})$ (top-left), $B(\sigma, \rho, \eta, ms_{fiobt}, ms_{iobt})$ (top-right), $A(.) - B(.)$ (bottom-left) and $A(.) / B(.)$ (bottom-right) varying ms_{fiobt} and ms_{iobt} while fixing $\rho = 4.0$, $\sigma = 2.5$ and $\eta = 1.2$. Each colored square indicates the value of the corresponding function (e.g., $A(.)$ for the top-left panel) for the pair of initial market shares (i.e., ms_{fiobt} and ms_{iobt}). The coloured cells of the top two figures and the bottom-left figure indicates the corresponding markup adjustments in percentages. The numbers in the coloured cells of the bottom-right figure give the ratio of the two reallocation effects (i.e., $A(.) / B(.)$) and are based on a different colour scheme than the other three figures due to the different scales.

C Supplementary estimation results

C.1 Origin-destination-year fixed effects

Table C10 presents results for specifications which include origin-destination-year fixed effects δ_{odt} instead of an indicator for trade agreements PTA_{odt} and origin-destination fixed effects δ_{od} . The results are qualitatively similar to but larger in magnitude than those in table 2.

Table C10: Elasticities to tariffs using origin-destination-year fixed effects

	Firm's mkt share in the dest. $\ln(\omega_{fiotd})$ (1)	Markups $\ln(\mu_{fiotd})$ (2)	Firm's within origin mkt share $\ln(ms_{fiotd})$ (3)	Origin's mkt share $\ln(ms_{iotd})$ (4)	No. of firms (PPML) (5)
Tariff $_{iotd}$	-1.38*** (0.302)	0.42*** (0.088)	3.58*** (0.402)	-3.15*** (0.347)	-2.72*** (0.196)
R ²	0.66	0.90	0.79	0.78	-
Observations	15,852,024	15,791,798	15,852,024	1,064,761	2,728,200
Fixed Effects					
Firm-product-origin-year	✓	✓	✓		
Product-origin-year				✓	✓
Product-destination-year	✓	✓	✓	✓	✓
Origin-destination-year	✓	✓	✓	✓	✓

Notes: The dependent variable is the firm's log (tariff-inclusive) export value in column (1), the firm's log (tariff-exclusive) unit value in column (2), the log of the firm's share of its country's trade with the destination in column (3), the log of the country's (tariff-inclusive) export value to the destination market in column (4) and the number of firms in column (5). Standard errors, reported in parentheses, are clustered at the product-destination level, and we denote statistical significance with *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$. Estimates are based on an integrated dataset of firms' exports from eleven countries built from the World Bank Exporter Dynamics Database, China's Customs Authority, and Egypt's Customs Authority, as well as tariff data from the WTO and Feenstra and Romalis (2014), and the World Bank Deep Trade Agreements Database.

C.2 Extreme tariffs

Since some tariffs in our dataset are extremely large, table C11 presents results for a subset of our data which excludes all observations where the destination applies a tariff in excess of 40%, the 99th percentile of the tariff distribution in our sample, on that product against any of its trading partners in that year. The results are again qualitatively similar to but larger in magnitude than those in table 2.

Table C11: Elasticities excluding tariff outliers

	Firm's mkt share in the dest. $\ln(\omega_{fiotd})$ (1)	Markups $\ln(\mu_{fiotd})$ (2)	Firm's within origin mkt share $\ln(ms_{fiotd})$ (3)	Origin's mkt share $\ln(ms_{ioidt})$ (4)	No. of firms (PPML) (5)
Tariff _{ioidt}	-0.91*** (0.264)	0.46*** (0.080)	3.45*** (0.326)	-3.01*** (0.255)	-2.45*** (0.172)
PTA _{oidt}	0.01 (0.022)	-0.02*** (0.009)	0.01 (0.029)	-0.05* (0.023)	0.00 (0.009)
R ²	0.65	0.90	0.79	0.77	-
Observations	15,591,039	15,531,010	15,591,039	1,059,219	2,718,679
Fixed Effects					
Firm-product-origin-year	✓	✓	✓		
Product-origin-year				✓	✓
Product-destination-year	✓	✓	✓	✓	✓
Origin-destination	✓	✓	✓	✓	✓

Notes: The dependent variable is the firm's log (tariff-inclusive) export value in column (1), the firm's log (tariff-exclusive) unit value in column (2), the log of the firm's share of its country's trade with the destination in column (3), the log of the country's (tariff-inclusive) export value to the destination market in column (4) and the number of firms in column (5). Standard errors, reported in parentheses, are clustered at the product-destination level, and we denote statistical significance with *** p<0.01, ** p<0.05, and * p<0.1. Estimates are based on an integrated dataset of firms' exports from eleven countries built from the World Bank Exporter Dynamics Database, China's Customs Authority, and Egypt's Customs Authority, as well as tariff data from the WTO and Feenstra and Romalis (2014), and the World Bank Deep Trade Agreements Database.

C.3 Elasticities with standard errors clustered at the firm level

Table C12 presents results with standard errors clustered at the firm level for the three dependent variables which vary at the firm level. While the standard errors on the tariff coefficient in the regressions for the firm’s market share in the destination $\ln(\omega_{fi\text{odt}})$ and the firm’s markup $\ln(\mu_{fi\text{odt}})$ are higher than in table 2, the opposite is true for the standard error on the tariff coefficient in the regression for the firm’s within origin market share $\ln(ms_{fi\text{odt}})$.

Table C12: Elasticities with standard errors clustered at the firm level

	Firm’s mkt share in the dest. $\ln(\omega_{fi\text{odt}})$ (1)	Markups $\ln(\mu_{fi\text{odt}})$ (2)	Firm’s within origin mkt share $\ln(ms_{fi\text{odt}})$ (3)	Origin’s mkt share $\ln(ms_{i\text{odt}})$ (4)	No. of firms (PPML) (5)
Tariff _{i\text{odt}}}	-0.78* (0.451)	0.41** (0.169)	2.86*** (0.240)	-2.31*** (0.271)	-2.21*** (0.162)
PTA _{odt}	0.01 (0.043)	-0.02 (0.014)	0.05* (0.027)	-0.04* (0.023)	-0.05*** (0.009)
R ²	0.66	0.90	0.79	0.77	-
Observations	15,853,618	15,793,386	15,853,618	1,067,240	2,750,833
Fixed Effects					
Firm-product-origin-year	✓	✓	✓		
Product-origin-year				✓	✓
Product-destination-year	✓	✓	✓	✓	✓
Origin-destination	✓	✓	✓	✓	✓

Notes: The dependent variable is the firm’s log (tariff-inclusive) export value in column (1), the firm’s log (tariff-exclusive) unit value in column (2), the log of the firm’s share of its country’s trade with the destination in column (3), the log of the country’s (tariff-inclusive) export value to the destination market in column (4) and the number of firms in column (5). Standard errors, reported in parentheses, are clustered firm level in columns (1) - (3) and the product-destination level in columns (4) and (5), and we denote statistical significance with *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$. Estimates are based on an integrated dataset of firms’ exports from eleven countries built from the World Bank Exporter Dynamics Database, China’s Customs Authority, and Egypt’s Customs Authority, as well as tariff data from the WTO and Feenstra and Romalis (2014), and the World Bank Deep Trade Agreements Database.