



HAL
open science

Product Mix and Firm Productivity Responses to Trade Competition

Thierry Mayer, Marc Melitz, Gianmarco Ottaviano

► **To cite this version:**

Thierry Mayer, Marc Melitz, Gianmarco Ottaviano. Product Mix and Firm Productivity Responses to Trade Competition. 2016. hal-03393193

HAL Id: hal-03393193

<https://sciencespo.hal.science/hal-03393193>

Preprint submitted on 21 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

DISCUSSION PAPER SERIES

DP11389

**PRODUCT MIX AND FIRM
PRODUCTIVITY RESPONSES TO TRADE
COMPETITION**

Thierry Mayer, Marc J Melitz and Gianmarco
Ottaviano

***INTERNATIONAL TRADE AND
REGIONAL ECONOMICS***



PRODUCT MIX AND FIRM PRODUCTIVITY RESPONSES TO TRADE COMPETITION

Thierry Mayer, Marc J Melitz and Gianmarco Ottaviano

Discussion Paper 11389

Published 14 July 2016

Submitted 14 July 2016

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INTERNATIONAL TRADE AND REGIONAL ECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Thierry Mayer, Marc J Melitz and Gianmarco Ottaviano

PRODUCT MIX AND FIRM PRODUCTIVITY RESPONSES TO TRADE COMPETITION

Abstract

We document how demand shocks in export markets lead French multi-product exporters to re-allocate the mix of products sold in those destinations. In response to positive demand shocks, those French firms skew their export sales towards their best performing products; and also extend the range of products sold to that market. We develop a theoretical model of multi-product firms and derive the specific demand and cost conditions needed to generate these product-mix reallocations. Our theoretical model highlights how the increased competition from demand shocks in export markets – and the induced product mix reallocations – induce productivity changes within the firm. We then empirically test for this connection between the demand shocks and the productivity of multi-product firms exporting to those destinations. We find that the effect of those demand shocks on productivity are substantial – and explain an important share of aggregate productivity fluctuations for French manufacturing.

JEL Classification: F1

Keywords:

Thierry Mayer - thierry.mayer@sciences-po.fr
Sciences Po, Banque de France, CEPII and CEPR

Marc J Melitz - mmelitz@harvard.edu
Harvard University and CEPR

Gianmarco Ottaviano - g.i.ottaviano@lse.ac.uk
LSE and CEPR

Acknowledgements

This research has received funding from the European Research Council (ERC) under the Grant Agreement No. 313522. For many helpful comments and suggestions we thank Pol Antras, Nick Bloom, Paola Conconi, Gita Gopinath, Elhanan Helpman, Eduardo Morales, Matthieu Parenti, John Van Reenen, Jean-Marc Robin, Dan Trefler; as well our discussants Mary Amiti, Peter Schott, and Volker Nocke; and participants at many seminars and conferences.

Product Mix and Firm Productivity Responses to Trade Competition*

Thierry Mayer
Sciences-Po, Banque de France
CEPII and CEPR

Marc J. Melitz
Harvard University
NBER and CEPR

Gianmarco I.P. Ottaviano
London School of Economics
U Bologna, CEP and CEPR

June 2016

Abstract

We document how demand shocks in export markets lead French multi-product exporters to re-allocate the mix of products sold in those destinations. In response to positive demand shocks, those French firms skew their export sales towards their best performing products; and also extend the range of products sold to that market. We develop a theoretical model of multi-product firms and derive the specific demand and cost conditions needed to generate these product-mix reallocations. Our theoretical model highlights how the increased competition from demand shocks in export markets – and the induced product mix reallocations – induce productivity changes within the firm. We then empirically test for this connection between the demand shocks and the productivity of multi-product firms exporting to those destinations. We find that the effect of those demand shocks on productivity are substantial – and explain an important share of aggregate productivity fluctuations for French manufacturing.

*This research has received funding from the European Research Council (ERC) under the Grant Agreement No. 313522. For many helpful comments and suggestions we thank Pol Antras, Nick Bloom, Paola Conconi, Gita Gopinath, Elhanan Helpman, Eduardo Morales, Matthieu Parenti, John VanReenen, Jean-Marc Robin, Dan Treﬂer; as well our discussants Mary Amiti, Peter Schott, and Volker Nocke; and participants at many seminars and conferences.

1 Introduction

Recent studies using detailed micro-level datasets on firms, plants, and the products they produce have documented vast differences in all measurable performance metrics across those different units. Those studies have also documented that these performance differences are systematically related to participation in international markets (see, e.g., Bernard, Jensen, Redding and Schott (2007, 2012) for the U.S., and Mayer and Ottaviano (2008) for Europe): Exporting firms and plants are bigger, more productive, more profitable, and less likely to exit than non-exporters. And better performing firms and plants export a larger number of products to a larger number of destinations. Exporters are larger in terms of employment, output, revenue and profit. Similar patterns also emerge across the set of products sold by multi-product firms. There is a stable performance ranking for firms based on the products' performance in any given market, or in worldwide sales. Thus, better performing products in one market are most likely to be the better performing products in any other market (including the global export market). This also applies to the products' selection into a destination, so better performing products are also sold in a larger set of destinations.

Given this heterogeneity, trade shocks induce many different reallocations across firms and products. Some of these reallocations are driven by 'selection effects' that determine which products are sold where (across domestic and export markets), along with firm entry/exit decisions (into/out of any given export market, or overall entry/exit of the firm). Other reallocations are driven by 'skewness effects' whereby – conditional on selection (a given set of products sold in a given market) – trade affects the relative market shares of those products. Both types of reallocations generate (endogenous) productivity changes that are independent of 'technology' (the production function at the product-level). This creates an additional channel for the aggregate gains from trade.

Unfortunately, measuring the direct impact of trade on those reallocations across firms is a very hard task. On one hand, shocks that affect trade are also likely to affect the distribution of market shares across firms. On the other hand, changes in market shares across firms likely reflect many technological factors (not related to reallocations). Looking at reallocations across products within firms obviates many of these problems. Recent theoretical models of multi-product firms highlight how trade induces a similar pattern of reallocations within firms as it does across firms. And measuring reallocations within multi-product firms has several advantages: Trade shocks that are exogenous to individual firms can be identified much more easily than at a higher level of aggregation; Controls for any technology changes at the firm-level are also possible; and

reallocations can be measured for the same set of narrowly defined products sold by the same firm across destinations or over time. In addition, impediments to factor reallocations are likely to be substantially higher across firms than across product lines within firms. Moreover, multi-product firms dominate world production and trade flows. Hence, reallocations within multi-product firms have the potential to generate large changes in aggregate productivity. Empirically, we find very strong evidence for the effects of trade shocks on those reallocations, and ultimately on the productivity of multi-product firms. The overall impact on aggregate French manufacturing productivity is substantial.

The rest of the paper is organized as follows. Section 2 briefly surveys the recent empirical literature on trade-induced reallocations. Section 3 introduces our dataset on French exporters and provides novel evidence on such reallocations with a special emphasis on skewness effects. It shows that positive demand shocks in any given destination market induce French exporters to skew their product level export sales to that destination towards their best performing products. These demand shocks also lead to strong positive responses in both the intensive and extensive margins of export sales to that destination. Section 4 introduces a flexible theoretical framework with multi-product firms to rationalize all this empirical evidence. Our framework highlights how certain properties of demand – which relate to the curvature of the demand and marginal revenue curves – are crucial in generating predictions for reallocations that are consistent with the data. In particular, these empirically relevant properties rule out the case of constant elasticity demand. They also imply that positive demand shocks engender increases in multi-product firm productivity since those shocks induce firms to produce relatively more of their better performing products (the skewness effect). Sections 5 and 6 take these predictions to the data and measure large responses in multi-product firm productivity to demand shocks in export markets. We find that these productivity responses explain an important share of aggregate productivity fluctuations for French manufacturing.

2 Previous Evidence on Trade-Induced Reallocations

In a previous paper, Mayer, Melitz and Ottaviano (2014), we investigated, both theoretically and empirically, the mechanics of product reallocations within multi-product firms across export destination markets. We used comprehensive firm-level data on annual shipments by all French exporters to all countries in the world (not including the French domestic market) for a set of more than 10,000 goods. Firm-level exports are collected by French customs and include export sales for each

8-digit (combined nomenclature or NC8) product by destination country. Our focus then was on the cross-section of firm-product exports across destinations (for a single year, 2003). We presented evidence that French multi-product firms indeed exhibit a stable ranking of products in terms of their shares of export sales across export destinations with ‘core’ products being sold in a larger number of destinations (and commanding larger market shares across destinations). We used the term ‘skewness’ to refer to the concentration of these export market shares in any destination and showed that this skewness consistently varied with destination characteristics such as GDP and geography: French firms sold relatively more of their best performing products in bigger, more centrally-located destinations (where competition from other exporters and domestic producers is tougher).

Other research has also documented similar patterns of product reallocations (within multi-product firms) over time following trade liberalization. For the case of CUSFTA/NAFTA, Baldwin and Gu (2009), Bernard, Redding and Schott (2011), and Iacovone and Javorcik (2010) all report that Canadian, U.S., and Mexican multi-product firms reduced the number of products they produce during these trade-liberalization episodes. Baldwin and Gu (2009) and Bernard, Redding and Schott (2011) further report that CUSFTA induced a significant increase in the skewness of production across products. Iacovone and Javorcik (2010) separately measure the skewness of Mexican firms’ export sales to the US. They report an increase in this skewness following NAFTA: They show that Mexican firms expanded their exports of their better performing products (higher market shares) significantly more than those for their worse performing exported products during the period of trade expansion from 1994 – 2003.

As prices are rarely observed, there is little direct evidence on how markups, prices and costs are related across products supplied with different productivity, and how they respond to trade liberalization.¹ A notable exception is the recent paper by De Loecker, Goldberg, Pavcnik and Khandelwal (2016) who exploit unique information on the prices and quantities of Indian firms’ products over India’s trade liberalization period from 1989 – 2003. They also document that better performing firms (higher sales and productivity) produce more products, before focusing on the evidence for firm and product level markups. Across firms, they show that those better performing firms also set higher markups. They also document a similar pattern across the products produced by a given multi-product firm, which sets relatively higher markups on their better performing

¹Prices are typically backed-out as unit values based on reported quantity information, which is extremely noisy.

products (lower marginal cost and higher market shares).² In addition, they show strong evidence for endogenous markup adjustments via imperfect pass-through from products' marginal costs to their prices: Only a portion of marginal cost decreases are passed on to consumers in the form of lower prices, while the remaining portion goes to higher markups. This is consistent with recent firm-level evidence on exchange rate pass-through. Berman, Martin and Mayer (2012) analyze the heterogeneous reaction of French exporters to real exchange rate changes across destinations over a decade (1995 – 2005). They find that on average firms react to depreciation by increasing their markup. They also find that high-performance firms increase their markup significantly more – implying that the pass-through rate is significantly lower for better performing firms.

Berman, Martin and Mayer (2012) also find very strong evidence that this pass-through rate (the elasticity of price with respect to the exchange rate) is heterogeneous across firms: it sharply decreases with firm performance (better performing firms adjust their markups substantially more than worse performing firms in response to the same exchange rate shock). Li, Ma and Xu (2015) confirm this result of decreasing pass-through (with firm performance) for Chinese exporters. Amiti, Itskoki, and Konings (2014) also find a very similar result for Belgian firms, whereby pass-through is strongly decreasing with the Belgian firms' export market share. Chatterjee, Dix-Carneiro, and Vichyanond (2013) also show this pattern for Brazilian exporters. In addition, they find that this property of decreasing pass-through also holds within multi-product Brazilian firms across their set of exported products. That is, they find substantially lower pass-through rates for a firm's better performing products (with relatively higher market shares).

This evidence on prices, markups, and pass-through is directly related to properties of demand (such as its elasticity). We review those connections when we introduce our theoretical model and its demand properties. We will then show that the same demand properties are required to generate the product reallocations that we document in the following section. This provides an independent source of confirmation for the evidence based on product-level prices.

3 Reallocations Over Time

We now document how changes *within* a destination market over time induce a similar pattern of reallocations as the ones we previously described (holding across destinations). More specifically, we show that demand shocks in any given destination market induce firms to skew their product

²Using product-level prices for multi-product firms, Dhyne et al (2016) also document a strong correlation between a product's sales rank and its productivity.

level export sales to that destination towards their best performing products. In terms of first moments, we show that these demand shocks also lead to strong positive responses in both the intensive and extensive margins of export sales to that destination.

3.1 Data

We use the same data source for firm exports (described in Section 2) as Mayer, Melitz and Ottaviano (2014), extended to cover multiple years from 1995-2005. All firms operating in the French metropolitan territory must report their export sales according to the following criteria: Exports to each EU destination whenever within-EU exports exceeds 100,000 Euros;³ and exports to non-EU country whenever exports to that destination exceeds 1,000 Euros or a ton. Despite these limitations, the database is nearly comprehensive. For instance, in 2005, 103,220 firms report exports across 234 destination countries (or territories) for 9873 products. This represents data on over 2.2 million shipments.

We restrict our analysis to firms whose main activity is classified as manufacturing to ensure that firms take part in the production of the goods they export.⁴ This leaves us with data covering more than a million shipments by firms across all manufacturing sectors.⁵

Matched balance-sheet data provide us with information on variables that are needed to assess firm productivity such as turnover, value added, employment, investment, raw material use and capital. However, we can only measure product reallocations in terms of sales in export markets as the breakdown of sales across products for the domestic market is not available to us. We will have to take this into account in designing our estimation strategy. The balance-sheet data we have access to comes in two sources where the official identification number of the firm can be matched with customs information. The first source is the EAE, produced by the national statistical institute, and exhaustive for manufacturing firms with size exceeding 20 employees. The second is BRN, which comes from tax authorities, and includes a broader coverage of firms, since it is based on the firm's legal tax regime and a relatively low sales threshold. Whenever firm data is

³If this threshold is not met, firms can choose to report under a simplified scheme without supplying separate EU export destinations. However, in practice, many firms under the threshold report their exports by destination. During our sample period, the threshold increased from 38,050 Euros (250,000 French Francs) to 100,000 Euros in 2001. We have checked that all our results are robust to using only the 1996-2001 subsample (a sample period that also features a more stable product classification schedule; See Pierce and Schott, 2012).

⁴Some large distributors such as Carrefour account for a disproportionate number of annual export shipments.

⁵In a robustness check, we also drop observations for firms that report an affiliate abroad. This avoids the issue that multinational firms may substitute exports of some of their best performing products with affiliate production in the destination country (affecting our measurement of product reallocations). Results are quantitatively very similar in all regressions.

available from both sources, we give precedence to the EAE data (which is more closely monitored by the statistical authorities).⁶ Table 1 provides some descriptive statistics relevant to the match between customs and balance sheet data. The overall match is not perfect but covers between 88 and 95 percent of the total value of French exports. The match with firms declaring manufacturing as their main activity is still very good although there is a clear trend of declining quality of match, particularly after 2000. This is also visible in the aggregate growth rate of exports in our sample (column 5) that overall provide a quite good match of the overall exports growth rate in column (4), but deteriorating over time. Our investigations suggest that the increasing propensity of large French manufacturers to declare their main activity as retail or some other service activity might provide part of that explanation.⁷ Overall our matched dataset is very comparable to recent papers using the same primary sources as in by Eaton et al. (2011), Berman et al. (2012) or di Giovanni et al. (2014) for instance.

Table 1: Aggregate Exports, Production, Employment, and Productivity

Year	Exports								
	Exports			VA	Emp.	VA/Emp.			
	Value	Share	Growth Rate			B Eur	M work.	K Eur	Growth
B Eur	Match	Mfg	Full	Mfg	B Eur	M work.	K Eur	Growth	
1995	211.3	94.9	74.0			177.1	2.903	61.0	
1996	219.6	93.5	73.2	3.9	2.8	178.6	2.918	61.2	0.3
1997	252.7	92.8	72.6	15.1	14.2	188.0	2.899	64.8	6.0
1998	267.1	92.0	72.1	5.7	4.9	193.3	2.914	66.3	2.3
1999	277.5	91.1	71.7	3.9	3.3	198.9	2.870	69.3	4.5
2000	319.4	90.8	71.9	15.1	15.4	209.5	2.924	71.6	3.4
2001	324.6	89.9	69.0	1.6	-2.3	199.4	2.932	68.0	-5.1
2002	321.7	90.4	68.4	-0.9	-1.8	198.0	2.865	69.1	1.7
2003	314.3	90.4	65.3	-2.3	-6.7	187.3	2.633	71.1	2.9
2004	335.0	88.4	64.6	6.6	5.4	193.2	2.577	75.0	5.4
2005	350.8	88.0	62.9	4.7	1.9	194.9	2.505	77.8	3.8

Column (1): Total value of exports from customs data. Column (2): Export share of firms matched with balance-sheet data from BRN/EAE. Column (3): Export share of matched firms with main activity classified as manufacturing. Columns (4) and (5): Export growth rates for customs data and subset of manufacturing firms. Columns (6) to (8): Aggregate value added, employment, and their ratio for the subset of manufacturing firms. Column (9): Growth rate of labor productivity from previous column.

⁶The correlations between variables reported both in EAE and BRN (value added, employment, exports, capital stock) are very high: between .91 and .99 in all cases.

⁷This is another reason why we check all our results against the early 1996-2001 sub-sample.

3.2 Measuring Export Demand Shocks

Consider a firm i exporting a number of products s in industry I to destination d in year t . We measure industries (I) at the 3-digit ISIC level (35 different classifications across French manufacturing). We consider several measures of demand shocks that affect this export flow. At the most aggregate level we use the variation in GDP in d , $\log \text{GDP}_{d,t}$. At the industry level I , we use total imports into d *excluding* French exports, $\log M_{d,t}^I$. We can also use our detailed product-level shipment data to construct a firm i -specific demand shock:

$$\text{trade shock}_{i,d,t}^I \equiv \overline{\log M_{d,t}^s} \quad \text{for all products } s \in I \text{ exported by firm } i \text{ to } d \text{ in year } t_0, \quad (1)$$

where $M_{d,t}^s$ represents total imports into d (again, excluding French exports) for product s and the overline represents the (unweighted) mean. For world trade, the finest level of product level of aggregation is the HS-6 level (from UN-COMTRADE and CEPII-BACI)⁸, which is more aggregated than our NC8 classification for French exports (roughly 5,300 HS products per year versus 10,000 NC8 products per year). The construction of the last trade shock is very similar to the one for the industry level imports $\log M_{d,t}^I$, except that we only use imports into d for the precise product categories that firm i exports to d .⁹ In order to ensure that this demand shock is exogenous to the firm, we use the set of products exported by the firm in its first export year in our sample (1995, or later if the firm starts exporting later on in our sample), and then exclude this year from our subsequent analysis. Note that we use an un-weighted average so that the shocks for all exported products s (within an industry I) are represented proportionately.¹⁰

For all of these demand shocks $X_t = \text{GDP}_{d,t}, M_{d,t}^I, M_{d,t}^s$, we compute the first difference as the Davis-Haltiwanger growth rate: $\tilde{\Delta}X_t \equiv (X_t - X_{t-1}) / (.5X_t + .5X_{t-1})$. This measure of the first difference preserves observations when the shock switches from 0 to a positive number, and has a maximum growth rate of -2 or 2 . This is mostly relevant for our measure of the firm-specific trade shock, where the product-level imports into d , $M_{d,t}^s$ can often switch between 0 and positive values. Whenever $X_{t-1}, X_t > 0$, $\tilde{\Delta}X_t$ is monotonic in $\Delta \log X_t$ and approximately linear for typical growth

⁸See Gaulier and Zignago (2010).

⁹There is a one-to-many matching between the NC8 and HS6 product classifications, so every NC8 product is assigned a unique HS6 classification. We use the same $M_{d,t}^s$ data for any NC8 product s within the same HS6 classification.

¹⁰Thus, positive idiosyncratic demand shocks to high market-share products (which mechanically contribute to increase the skewness of product sales) are given the same weight as positive idiosyncratic shocks to low market-share products (which mechanically contribute to decrease the skewness of product sales); and vice-versa for negative shocks.

rates ($|\Delta \log X_t| < 2$).¹¹ We thus obtain our three measures of demand shocks in first differences: $\tilde{\Delta} \text{GDP}_{d,t}$, $\tilde{\Delta} M_{d,t}^I$, $\overline{\tilde{\Delta} M_{d,t}^s}$. From here on out, we refer to these three shocks, respectively, as GDP shock, trade shock – ISIC, and trade shock. Note that this last firm-level shock, $\overline{\tilde{\Delta} M_{d,t}^s}$, represents the un-weighted average of the growth rates for all products exported by the firm in $t - 1$.

3.3 The Impact of Demand Shocks on Trade Margins and Skewness

Before focusing on the effects of the demand shocks on the skewness of export sales, we first show how the demand shocks affect firm export sales at the intensive and extensive margins (the first moments of the distribution of product export sales). Table 2 reports how our three demand shocks (in first differences) affect changes in firm exports to destination d in ISIC I (so each observation represents a firm-destination-ISIC combination). We decompose the firm’s export response to each shock into an intensive margin (average exports per product) and an extensive margin (number of exported products). We clearly see how all three demand shocks induce very strong (and highly significant) positive responses for both margins. This confirms that our demand shocks capture important changes in the local demand faced by French exporters.¹²

Table 2: Demand Shocks and Local Exports

Dependent Variable	$\Delta \log$ Exports per Product			$\Delta \log$ # Products Exported		
$\tilde{\Delta}$ GDP Shock	0.493 ^a (0.048)			0.149 ^a (0.016)		
$\tilde{\Delta}$ trade shock	0.277 ^a (0.009)			0.076 ^a (0.004)		
$\tilde{\Delta}$ trade shock - ISIC	0.039 ^a (0.005)			0.014 ^a (0.002)		
Observations	401575	407520	407520	401575	407520	407520

Standard errors in parentheses: ^c < 0.1, ^b < 0.05, ^a < 0.01. All regressions include year dummies, and standard errors clustered at the level relevant for the variable of interest: destination country for columns (1) and (4), firm-destination for columns (2) and (5) and ISIC-destination for columns (3) and (6).

¹¹Switching to first difference growth rates measured as $\Delta \log X_t$ (and dropping products with zero trade in the trade shock average) does not materially affect any of our results.

¹²Specifications using the log levels of the shocks and firm-destination-ISIC fixed-effects yield similar results. Other specifications including the three covariates show that those demand shocks are different enough to be estimated jointly while each keeping its positive sign and statistical significance. We also have experimented with the addition of a control for a worldwide shock (aggregated from the destination-specific shocks using the firm’s prior year export shares). The effect of this additional control on the destination-specific coefficient is negligible, thus confirming that worldwide shocks do not influence the coefficients presented in Table 2.

We now investigate the consequences of those demand shocks for the skewness of export sales (independent of the level of product sales). In Mayer, Melitz and Ottaviano (2014), we focused on those effects in the cross-section across destinations. Here, we examine the response of skewness within a destination over time using our new demand shocks. We rely on the Theil index as our measure of skewness due to its aggregation properties: We will later aggregate the export responses at the destination-*ISIC* level up to the firm-level – in order to generate predictions for firm-level productivity. Thus, our measure of skewness for the distribution of firm *i*'s exports to destination *d* in industry *I*, $x_{i,d,t}^s$, is the Theil index (computed over all N_{idt}^I products *s* that firm *i* exports to *d* in year *t*):

$$T_{i,d,t}^I \equiv \frac{1}{N_{idt}^I} \sum_{s \in I} \frac{x_{i,d,t}^s}{\bar{x}_{idt}^I} \log \left(\frac{x_{i,d,t}^s}{\bar{x}_{idt}^I} \right), \quad \bar{x}_{idt}^I \equiv \frac{\sum_{s \in I} x_{idt}^s}{N_{idt}^I}. \quad (2)$$

Table 3 reports regressions of this skewness measure on all three demand shocks (jointly) – at the firm-destination-*ISIC* level. In the first column, we use a specification in (log) levels (FE), and use firm-destination-*ISIC* fixed effects to isolate the variation over time. In the second column, we return to our specification in first differences (FD). In the third column we add the firm-destination-*ISIC* fixed effects to this specification in first differences (FD-FE). This controls for any trend growth rate in our demand shocks over time. Across all three specifications, we see that positive demand shocks induce a highly significant increase in the skewness of firm export sales to a destination. The effect of all three shocks are weakened a little bit due to some collinearity, in particular in the most-demanding FD-FE specification. To highlight this effect, we re-run our FD-FE specification with each shock entered independently. Those results are reported in the first column of Table 4. Each row is the result of that independent regression with each shock. In this case, even the weakest *ISIC*-level shock remains significant at the 1% level.¹³

The three remaining columns of Table 4 show additional robustness specifications with an alternative measure of skewness—the Atkinson index. This index was developed to allow for greater flexibility in quantifying the contribution of different parts of the distribution to overall inequality. It is defined as:

$$A_{i,d,t}^{I,\eta} \equiv \begin{cases} 1 - \frac{1}{\bar{x}_{idt}^I} \left[\frac{1}{N_{idt}^I} \sum_{s \in I} (x_{idt}^s)^{1-\eta} \right]^{\frac{1}{1-\eta}}, & \bar{x}_{idt}^I \equiv \frac{\sum_{s \in I} x_{idt}^s}{N_{idt}^I}, \quad \text{for } 0 \leq \eta \neq 1 \\ 1 - \frac{1}{\bar{x}_{idt}^I} \left[\prod_{s \in I} x_{idt}^s \right]^{\frac{1}{N_{idt}^I}}, & \text{for } \eta = 1 \end{cases} \quad (3)$$

¹³Just as we did for the impact of the destination-specific trade shocks on the first moment of exports (extensive and intensive margins), we experimented with the addition of a control for the firm's worldwide shock. Once again, the effect of this additional control on the destination-specific coefficient is negligible, thus confirming that worldwide shocks do not influence the coefficients presented in Table 4.

Table 3: Demand shocks and local skewness

Dependent Variable Specification	$T_{i,d,t}^I$	$\Delta T_{i,d,t}^I$	
	FE	FD	FD-FE
log GDP shock	0.075 ^a (0.016)		
log trade shock	0.048 ^a (0.005)		
log trade shock - ISIC	0.002 ^a (0.000)		
$\tilde{\Delta}$ GDP Shock		0.066 ^a (0.012)	0.066 ^a (0.015)
$\tilde{\Delta}$ trade shock		0.037 ^a (0.005)	0.033 ^a (0.006)
$\tilde{\Delta}$ trade shock - ISIC		0.006 ^a (0.002)	0.004 (0.003)
Observations	479387	401575	401575

Standard errors in parentheses: ^c < 0.1, ^b < 0.05, ^a < 0.01. FE refers to firm-destination-ISIC fixed effects. All regressions include year dummies, and standard errors clustered at the level of the destination country.

The parameter η is often called the inequality aversion parameter as higher values put more weight on the low end of the distribution. When η gets close to 0, more weight is given to high values, in our case to the firm's best performing products. As η increases, more weight is given to the distribution of the firm's worse performing products relative to the best performing ones. This Atkinson index is also equal to one minus the ratio of the generalized (CES) mean over the arithmetic mean. Different values of η provide different special cases of the generalized mean: with $\eta = 2$, it is the harmonic mean, whereas with $\eta = 1$ it is the geometric mean. In the limit, when $\eta = 0$, it is the arithmetic mean, in which case there can never be any inequality ($A_{i,d,t}^{I,\eta} \equiv 0$ for any distribution of product sales x_{idt}^s). Table 4 reproduces our skewness results for a wide range of inequality parameters η , starting with a low value of $\eta = .5$ along with the special cases of $\eta = 1$ and $\eta = 2$.¹⁴ Those results, reported in columns (2-4), clearly show that the effect of the demand shocks on skewness do not rely on the Theil functional form nor on a specific weighing scheme for the Atkinson index.

The skewness measures in Table 4 combine changes to both the intensive product mix margin (the relative sales of different products exported in both $t - 1$ and t), and the extensive margin

¹⁴This range of parameters is also consistent with that advocated by Atkinson and Brandolini (2010) for the evaluation of income inequality.

Table 4: Skewness and Shocks: Robustness - Skewness Measure

RHS	$\Delta T_{i,d,t}^I$	$\Delta A_{i,d,t}^{I,5}$	$\Delta A_{i,d,t}^{I,1}$	$\Delta A_{i,d,t}^{I,2}$
$\tilde{\Delta}$ GDP Shock	0.090 ^a (0.010)	0.039 ^a (0.004)	0.058 ^a (0.006)	0.058 ^a (0.007)
$\tilde{\Delta}$ Trade Shock	0.042 ^a (0.004)	0.018 ^a (0.002)	0.024 ^a (0.003)	0.023 ^a (0.003)
$\tilde{\Delta}$ Trade Shock - ISIC	0.008 ^a (0.002)	0.004 ^a (0.001)	0.006 ^a (0.001)	0.005 ^a (0.001)

Notes: This table presents the results of individual FD-FE regressions where the columns give the LHS and the rows give the RHS variables. All regressions include year and firm-destination-ISIC fixed effects. Standard errors in parentheses are clustered at the level of the destination country.

Table 5: Skewness and Shocks: Robustness - Intensive Margin

RHS	$\Delta T_{i,d,t}^{I,const}$	$\Delta A_{i,d,t}^{I,5,const}$	$\Delta A_{i,d,t}^{I,1,const}$	$\Delta A_{i,d,t}^{I,2,const}$	$\Delta \log R_{i,d,t}^I$
$\tilde{\Delta}$ GDP Shock	0.006 (0.008)	0.005 (0.004)	0.012 ^b (0.006)	0.014 ^b (0.007)	0.157 ^a (0.067)
$\tilde{\Delta}$ Trade Shock	0.012 ^a (0.003)	0.007 ^a (0.001)	0.011 ^a (0.003)	0.010 ^a (0.003)	0.059 ^b (0.027)
$\tilde{\Delta}$ Trade Shock - ISIC	0.005 ^a (0.002)	0.002 ^a (0.001)	0.004 ^a (0.001)	0.004 ^a (0.001)	0.008 (0.013)

Notes: This table presents the results of individual FD-FE regressions where the columns give the LHS and the rows give the RHS variables. All regressions include year and firm-destination-ISIC fixed effects. Standard errors in parentheses are clustered at the level of the destination country.

(products sold in one period but not the other). We now isolate the skewness response that is driven only by the intensive product mix margin. To do this, we recompute the skewness measures from Table 4 on the restricted subset of products that are exported in both $t - 1$ and t (denoted with a “const” superscript). Results using these intensive margin skewness measures are reported in the first 4 columns of Table 5. In column 5, we introduce one last skewness measure for the intensive margin: the sales ratio (in logs) of the firm’s best to second best performing product (in terms of global export sales). The effect of our firm-level trade shock is significant across all these specifications (and significant well beyond the 1% level for all the distribution wide skewness measures). The intensive margin skewness results for the ISIC-level trade shock and the GDP shock are always positive and significant in most (but not all) of those specifications.

4 Theoretical Framework

In the previous section we documented the pattern of product reallocations in response to demand shocks in export markets. We now develop a theoretical model of multi-product firms that highlights the specific demand conditions needed to generate this pattern. We show that these demand conditions are consistent with all of the micro-level evidence on firm/product selections, prices, and markups presented in section 2. In particular, those demand conditions highlight how demand shocks lead to changes in competition for exporters in those markets (which lead to the observed reallocations). Those reallocations, in turn, generate changes in firm productivity. We directly investigate the empirical connection between demand shocks and productivity in the following sections.

Our theoretical model contributes to the growing literature emphasizing demand systems with variable price elasticities for models of trade.¹⁵ Our main point of departure relative to those papers is that we seek to connect the demand conditions on variable elasticities (and further properties of demand) directly to the evidence on the product reallocations that we document; and our additional goal of empirically connecting those to firm-level productivity. In this spirit, our paper is most closely related to Arkolakis et al (2015) and Asplund and Nocke (2006), who also directly connect their theoretical modeling of endogenous markups with empirical moments: Arkolakis et al (2015) focus on the implications for the welfare gains from trade; and Asplund and Nocke (2006) focus on the implications for firm turnover in a dynamic setting.

¹⁵See, for example, Bertolotti and Epifani (2014), Dhingra and Morrow (2014), Fabinger and Weyl (2014), Mrazova and Neary (2014), and Zhelebodko et al (2012).

4.1 Closed Economy

To better highlight the role played by the properties of the demand system, we initially start with a closed economy and analyze the impact of demand shocks on product reallocations. We then develop an open economy version of our model that meshes more closely with our empirical setup. We show how demand shocks in an export market induce very similar product reallocations as in the closed economy (but for exporters to that market). We develop both a general equilibrium (with a single differentiated good sector for the whole economy) and a partial equilibrium version focusing on a single sector among many in the economy. In the latter, we also introduce a short-run version where entry is restricted general equilibrium is inherently a long-run scenario). We show how demand shocks induce the same skewness pattern in all of these modeling alternatives. This highlights the critical role of the demand system in shaping the pattern of reallocations.

Multi-Product Production with Additive Separable Utility

We consider a sector with a single productive factor, labor. We will distinguish between two scenarios. The first is the standard general equilibrium (GE) setup with a single sector. The exogenous labor endowment L indexes both the number of workers L^w (with inelastic supply) and consumers L^c . We choose the endogenous wage as the *numeraire*. Aggregate expenditures are then given by the exogenous labor endowment. In our partial equilibrium (PE) scenario, we focus on the sector as a small part of the economy. We take the number of consumers L^c as well as their individual expenditures on the sector's output as exogenously given. The supply of labor L^w to the sector is perfectly elastic at an exogenous economy-wide wage. We choose units so that both this wage and the exogenous expenditure per-consumer is equal to 1. This involves a normalization for the measure of consumers L^c and the choice of labor as *numeraire*: Aggregate accounting then implies that this normalized number of consumers L^c represents a fraction of the labor endowment L .¹⁶ In both scenarios, we model a demand shock as an increase in the number of consumers L^c . This increases aggregate expenditures one-for-one, given our assumptions of unitary consumer income. In our GE scenario, this increase is associated with a proportional increase in

¹⁶As we will restrict our analysis to additively separable preferences – which are non-homothetic – changes in consumer income will have different effects than changes in the number of consumers L^c . We focus on this functional form for tractability and do not wish to emphasize its properties for income elasticities. As first highlighted by Deaton and Muellbauer (1980), additively separable preference imply a specific relationship between price and income elasticities. We emphasize the properties of demand for those price elasticities. Thus, we analyze changes in the number of consumers L^c holding their income fixed. This is akin to indexing the preferences to a given reference income level.

labor supplied L^w . In our PE scenario, the labor supply response is left unrestricted, isolating the demand-side effects.

In both scenarios each consumer's utility is assumed to be additively separable over a continuum of imperfectly substitutable products indexed by $i \in [0, M]$ where M is the measure of products available. The representative consumer then solves the following utility maximization problem:

$$\max_{q_i \geq 0} \int_0^M u(q_i) di \text{ s.t. } \int_0^M p_i q_i di = 1,$$

where $u(q_i)$ is the sub-utility associated with the consumption of q_i units of product i . We assume that this sub-utility exhibits the following properties:

$$\mathbf{(A1)} \quad u(q_i) \geq 0 \text{ with equality for } q_i = 0; \quad u'(q_i) > 0 \text{ and } u''(q_i) < 0 \text{ for } q_i \geq 0.$$

The first order condition for the consumer's problem determines the *inverse* residual demand function (per consumer):

$$p(q_i) = \frac{u'(q_i)}{\lambda}, \tag{4}$$

where $\lambda = \int_0^M u'(q_i) q_i di > 0$ is the marginal utility of income. Given our assumption of separable preferences, this marginal utility of income λ is the unique endogenous aggregate demand shifter: Higher λ shifts all residual demand curves *inward*; we refer to this as an increase in competition for a given level of market demand L^c . Concavity of $u(q_i)$ ensures that the chosen consumption level from (4) also satisfies the second order condition for the consumer's problem. This residual demand curve (4) is associated with a marginal revenue curve

$$\phi(q_i) = \frac{u'(q_i) + u''(q_i) q_i}{\lambda}. \tag{5}$$

Let $\varepsilon_p(q_i) \equiv -p'(q_i) q_i / p(q_i)$ and $\varepsilon_\phi(q_i) \equiv -\phi'(q_i) q_i / \phi(q_i)$ denote the elasticities of inverse demand and marginal revenue (expressed in absolute values). Thus $\varepsilon_p(q_i) \geq 0$ is the inverse price elasticity of demand (less than 1 for elastic demand). Although the demand and marginal revenue curves are residual (they depend on λ), their elasticities are nonetheless independent of λ . These preferences nest the C.E.S. case where the elasticities $\varepsilon_p(q_i)$ and $\varepsilon_\phi(q_i)$ are constant; using the additively separable functional form for C.E.S., the marginal utility of income λ is then an inverse monotone function of the C.E.S. price index.

Products are supplied by firms that may be single- or multi-product. Market structure is

monopolistically competitive as in Mayer, Melitz and Ottaviano (2014): each product is supplied by a single firm and each firm supplies a countable number of products (among the continuum of consumed products). Technology exhibits increasing returns to scale associated with a fixed overhead cost, along with a constant marginal cost of production. We assume that the fixed cost f is common for all products while the marginal cost v (variety level cost) is heterogeneous. For a given firm, products are indexed in increasing order m of marginal cost from a ‘core product’ indexed by $m = 0$. Firm entry incurs a sunk cost f^e . After this cost is incurred, entrants randomly draw the marginal cost for their core product from a common continuous differentiable distribution $\Gamma(c)$ with support over $[0, \infty)$; we refer to this cost draw as the firm’s core competency.¹⁷ This gives an entrant the exclusive blueprint used to produce a countable range of additional products indexed by the integer m (potentially zero) at marginal cost $v(m, c) \equiv cz(m)$ with $z(0) = 1$ and $z'(m) > 0$.¹⁸

Product-Level Performance and Selection

A firm owning the blueprint for product i with marginal cost v and facing demand conditions λ chooses the optimal output per consumer $q(v, \lambda)$ to maximize total product-level profits $L^c [p(q_i)q_i - vq_i] - f$, so long as those profits are non-negative; or does not produce product i otherwise.¹⁹ The first order condition whenever production occurs equalizes marginal revenue with marginal cost:

$$\phi(q(v, \lambda)) = v. \tag{6}$$

In order to ensure that the solution to this maximization problem exists (for at least some $v > 0$) and is unique, we further restrict our choice of preferences to satisfy:

$$\mathbf{(A2)} \quad \varepsilon_p(0) < 1 \quad \text{and} \quad \mathbf{(A3)} \quad \varepsilon_\phi(q_i) > 0 \text{ for } q_i \geq 0.$$

These assumptions ensure that marginal revenue is decreasing for all output levels (A3) and positive (elastic demand) for at least some output levels (A2).

¹⁷This assumption of infinite support is made for simplicity in order to rule out the possibility of an equilibrium without any firm selection. We could also introduce an upper-bound cost draw so long as this upper-bound is high enough that the equilibrium features selection (some firms do not produce).

¹⁸The assumption $z'(m) > 0$ will generate the within-firm ranking of products discussed in Section 2. In the limit case when $z'(m)$ is infinite, all firms only produce a single product.

¹⁹As we have assumed that any firm’s set of product is of measure zero relative to the set of available products M , there is no product inter-dependence in the firm’s pricing/output decision (no cannibalization).

The profit maximizing price associated with the output choice (6) can be written in terms of the chosen markup $\mu(q_i) \equiv 1/(1 - \varepsilon_p(q_i))$:

$$p(q(v, \lambda)) = \mu(q(v, \lambda))v. \quad (7)$$

Those output and price choices are associated with the following product-level revenues and operating profits *per-consumer*:

$$r(v, \lambda) = p(q(v, \lambda))q(v, \lambda) \quad \text{and} \quad \pi(v, \lambda) = [p(q(v, \lambda)) - v]q(v, \lambda).$$

Total product-level sales are then given by $L^c r(v, \lambda)$ while total product net profits are $L^c \pi(v, \lambda) - f$. Using the first order condition for profit maximization (6) and our derivations for marginal revenue (5) and markup (7), the elasticities for all these product-level performance measures can be written in term of the elasticities of demand and marginal revenue:

$$\varepsilon_{q,v} = -\frac{1}{\varepsilon_\phi}, \quad \varepsilon_{q,\lambda} = -\frac{1}{\varepsilon_\phi} \quad (8)$$

$$\varepsilon_{r,v} = -\frac{1 - \varepsilon_p}{\varepsilon_\phi}, \quad \varepsilon_{r,\lambda} = -\frac{1 - \varepsilon_p}{\varepsilon_\phi} - 1, \quad (9)$$

$$\varepsilon_{\pi,v} = -\frac{1 - \varepsilon_p}{\varepsilon_p}, \quad \varepsilon_{\pi,\lambda} = -\frac{1}{\varepsilon_p}, \quad (10)$$

where we use the elasticity notation $\varepsilon_{g,x}$ to denote the elasticity of the function g with respect to x . All of the elasticities with respect to the product marginal cost v are negative, indicating that lower marginal cost is associated with higher output, sales, and profit (both operating and net). As expected, an increase in competition λ (for any given level of market demand L^c) will result in lower output, sales, and profit for all products.

Since operating profit is monotonic in a product's marginal cost v , the production decision associated with non-negative net profit will lead to a unique cutoff cost level \hat{v} satisfying

$$\pi(\hat{v}, \lambda)L^c = f. \quad (11)$$

All products with cost $v \leq \hat{v}$ will be produced. Since $v(0, c) = c$ (recall that $v(m, c) = cz(m)$; $z(0) = 1$) any firm with core competency $c \leq \hat{v}$ will produce at least its core ($m = 0$) product. Thus, $\hat{c} = \hat{v}$ will be the firm-level survival cutoff. Those surviving firms (with $c \leq \hat{c}$) will produce

$M(c) = \max \{m \mid cz(m) \leq \hat{v}\}$ *additional* products (potentially none, and then $M(c) = 0$) and earn firm-level net profits:

$$\Pi(c, \lambda) = \sum_{m=0}^{M(c)} [\pi(cz(m), \lambda) L^c - f].$$

Since $\varepsilon_{\pi, v}$ and $\varepsilon_{\pi, \lambda}$ are both negative, increases in competition λ – holding market demand L^c constant – will be associated with lower cutoffs $\hat{v} = \hat{c}$. Tougher competition thus leads to tougher selection: the least productive firms exit and all surviving firms shed (weakly) their worst performing products.

Free Entry in the Long-Run

In the long-run when entry is unrestricted, the expected profit of the prospective entrants adjusts to match the sunk cost:

$$\int_0^{\hat{c}} \Pi(c, \lambda) d\Gamma(c) = \sum_{m=0}^{\infty} \left\{ \int_0^{\hat{c}/z(m)} [\pi(cz(m), \lambda) L^c - f] d\Gamma(c) \right\} = f^e. \quad (12)$$

This free entry condition, along with the zero cutoff profit condition (11) jointly determine the equilibrium cutoffs $\hat{v} = \hat{c}$ along with the competition level λ . The number of entrants N^e is then determined by the consumer's budget constraint:

$$N^e \left\{ \sum_{m=0}^{\infty} \left[\int_0^{\hat{c}/z(m)} r(cz(m), \lambda) d\Gamma(c) \right] \right\} = 1. \quad (13)$$

These conditions hold in both our GE and PE scenarios.

Since labor is the unique factor and *numeraire*, we can convert the firms' costs (both per-unit production costs and the fixed costs) into employment. Aggregating over all firms yields the aggregate labor demanded:

$$L^w = N^e \left(f^e + \sum_{m=0}^{\infty} \left\{ \int_0^{\hat{c}/z(m)} [cz(m)q(cz(m), \lambda) L^c + f] d\Gamma(c) \right\} \right)$$

As the free entry condition (12) entails no ex-ante aggregate profits (aggregate revenue is equal to the payments to all workers, including those employed to cover the entry costs), this aggregate labor demand L^w will be equal to the number of consumers L^c . This ensures labor market clearing in our GE scenario. In our PE scenario, this implies that the endogenous labor supply adjusts so that it equalizes the normalized number of consumers (recall that this is an exogenous fraction of

the economy-wide labor endowment).²⁰

In these long-run scenarios, the impact of an increase in demand L^c on the cutoffs will depend on some further assumptions on demand, which we discuss in detail following the introduction of our short-run scenario. However, we note that the impact for the level of competition λ is unambiguous: higher demand leads to increases in competition λ .²¹

Short-Run Scenario

We now consider an alternative short-run situation in which the number of incumbents is fixed at \bar{N} in the PE scenario (with the same exogenous distribution of core competencies $\Gamma(c)$). In this case, free entry (12) no longer holds: firms with core competencies below the profit cutoff in (11) produce while the remaining firms shut-down. However, the budget constraint (13) still holds with the exogenous number of incumbents \bar{N} now replacing the endogenous number of entrants N^e . Together with the zero cutoff profit (11), those two conditions jointly determine the endogenous cutoffs $\hat{v} = \hat{c}$ and competition level λ .

As was the case in the long-run scenarios with free-entry, an increase in demand L^c must lead to an increase in competition λ .²² However, the response of the cutoffs is now unambiguous: the increase in demand must raise profits for all products (given their cost v) since there is no induced response in entry; leading to an increase in the cutoffs $\hat{v} = \hat{c}$.²³ In this short-run scenario, it is the production of these new varieties (previously unprofitable) that generates the increased competition. In the long run, the increased competition is driven by the entry of new firms (and the varieties they produce).

²⁰This sector-level adjustment for labor supply is very similar to the case of C.E.S. product differentiation within sectors and Cobb-Douglas preferences across sectors. In the latter case, the sector's labor supply adjusts so that it is equal (as a fraction of the aggregate labor endowment) to the exogenous Cobb-Douglas expenditure share for the sector. See Melitz & Redding (2014) for an example of those preferences with firm heterogeneity.

²¹Competition λ must strictly increase for the free entry condition to hold. If λ decreased (even weakly), then net profit for a given variety $\pi(v, \lambda)L^c - f$ should strictly increase (given $dL^c > 0$ and $\varepsilon_{\pi, \lambda} < 0$). The cutoffs $\hat{v} = \hat{c}$ must then strictly increase (given 11); along with the average firm profit (the left-hand-side) in the free-entry condition (12). Thus, this condition could not hold if λ weakly decreased.

²²Competition λ must strictly increase in order to satisfy the budget constraint (13) with a fixed number of incumbents. As we argued in a previous footnote (21), λ weakly decreasing would imply a strict increase in the cutoffs $\hat{v} = \hat{c}$. This would entail both a strict increase in the number of products consumed as well as a weak increase in expenditures per-product $r(v, \lambda)$ for each consumer (recall that $\varepsilon_{r, \lambda} < 0$). This would necessarily violate the consumer's budget constraint.

²³Note that with a fixed number of incumbents, the budget constraint (13) implies an increasing relationship between the cutoff \hat{c} and the level of competition λ (which reduces product-level revenues $r(v, \lambda)$ for all products).

Curvature of Demand

Up to now, we have placed very few restrictions on the shape of the residual demand curves that the firms face, other than the conditions (A1)-(A3) needed to ensure a unique monopolistic competition equilibrium. In particular, the rates of change of the elasticities of residual demand and marginal revenue (the signs of $\varepsilon'_p(q_i)$ and $\varepsilon'_\phi(q_i)$) were left unrestricted. We now show how further restrictions on those rates of change (or alternatively the curvature of demand and marginal revenue) are intrinsically tied to product-level reallocations – in addition to their better known consequences for prices, markups and pass-through. After we develop the open economy version of our model in the next section, we highlight the reverse connection from the pattern of product reallocations we documented in Section 3 back to their necessary conditions for demand. Throughout, we assume that conditions (A1)-(A3) hold in our monopolistic competition equilibrium. Thus, a *necessary* condition for an empirical pattern is an *additional* condition that must hold (to generate that empirical prediction) *conditional* on those initial assumptions.

The further restrictions on the shapes of demand and marginal revenue are:

$$\text{(B1)} \quad \varepsilon'_p(q_i) > 0 \text{ for } q_i \geq 0,$$

and

$$\text{(B1')} \quad \varepsilon'_\phi(q_i) > 0 \text{ for } q_i \geq 0.$$

Figure 1 depicts a log-log graph of the inverse demand and marginal revenue curves satisfying those restrictions. Under (B1) demand becomes more inelastic with consumption. This is also known as Marshall’s Second Law of Demand.²⁴ It is a necessary and sufficient condition for better performing products (lower v) to have higher markups and for tougher competition (higher λ) to lower markups for any given product (given a cost v).²⁵ Thus, the evidence discussed in Section

²⁴We thank Peter Neary for bringing this reference to our attention. (B1) is equivalent to the assumption by Arkolakis et al (2015) that the demand function of any product is log-concave in log-prices. In the terminology of Neary and Mrazova (2014) it defines the “sub-convex” case, with “sub-convexity” of an inverse demand function at an arbitrary point being equivalent to the function being less convex at that point than a CES demand function with the same elasticity. In the terminology of Zhelobodko et al (2012), this class of preferences exhibits increasing “relative love of variety” (RLV) as consumers care more about variety when their consumption level is higher. In the terminology of Kimball (1995), this class of preferences exhibits a positive superelasticity of demand. Bertoletti and Epifani (2014) refer to this case as “decreasing elasticity of substitution”. The “Adjustable pass-through” (APT) class of demand functions proposed by Fabinger and Weyl (2012) also satisfies (B1). Lastly, condition (B1) is sufficient for the induced maximized profit function to exhibit all the properties postulated by Asplund and Nocke (2006) in order to explain the empirical relationship between market size and firm turnover that they document (in a dynamic setting).

²⁵Since both higher cost v and higher λ is associated with lower output $q(v, \lambda)$; See (8).

2, linking better product and firm performance to higher markups implies that (B1) must hold. It is also consistent with the estimates of Arkolakis et al (2015) for bilateral trade demand. In our model with monopolistic competition, (B1) is also equivalent to an alternate condition that the pass-through elasticity from marginal cost to price $\theta \equiv \partial \ln p(q(v, \lambda)) / \partial \ln v = \varepsilon_p / \varepsilon_\phi$ is less than 1 (see appendix for proof). Thus, the vast empirical evidence on incomplete pass-through (see the survey by Burstein and Gopinath, 2014) also requires demand condition (B1). Most importantly, we will show that (B1) is a necessary condition for the evidence we documented on product reallocations for French exporters. This provides an independent confirmation for this demand condition that does not rely on the (very noisy) measurement of product prices and the estimation of markups (or alternatively marginal costs).

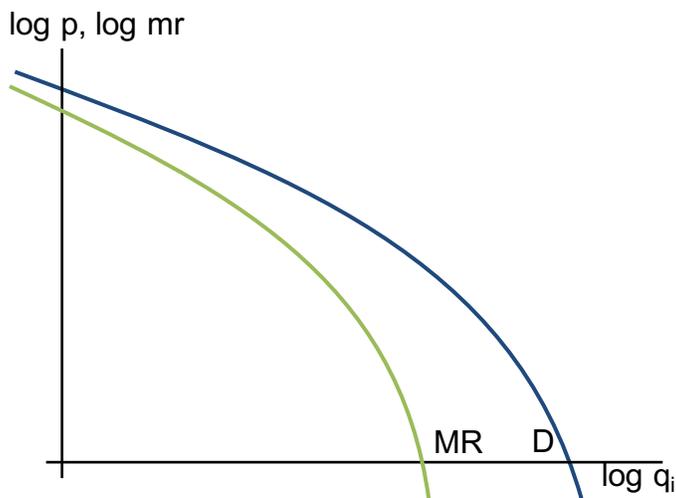


Figure 1: Graphical Representation of Demand Assumptions

Assumption (B1') is more restrictive than (B1): in the appendix, we show how (B1') implies (B1). If only (B1) holds, then the log demand curve would have the shape shown in Figure 1, but the log marginal revenue curve would need not be globally concave. However, it would still have to be everywhere steeper than log demand: $\varepsilon_\phi(q_i) > \varepsilon_p(q_i)$ for all $q_i \geq 0$ (see appendix). Whereas (B1) entails an increasing relationship between output and markup (and hence the *level* of pass-through), (B1') adds a connection between *changes* in output and *changes* in markups (and hence *differences* in pass-through). More specifically, (B1') holds if and only if the pass-through rate θ is increasing with marginal cost v : better performing products (lower v) have lower (more incomplete) pass-through: their markups adjust by more (in response to a given percentage shock to marginal cost) than worse performing products. Thus, (B1') is a necessary condition for

the evidence on such pass-through differences that we previously discussed (for French, Belgian, Brazilian, and Chinese firms/products). In the following section, we show that (B1') is a sufficient condition for the product reallocations we previously described for French exporters, in addition to all the evidence from the existing literature on prices, markups, and pass-through.²⁶

Lastly we note that conditions (B1) and (B1') exclude the C.E.S. case, where the derivative of the elasticities $\varepsilon'_p(q_i)$ and $\varepsilon'_\phi(q_i)$ are zero. In this limiting case, the inverse demand and marginal revenue in Figure 1 are linear. Nevertheless, (B1) and (B1') are consistent with most of the functional forms that have been used to explore endogenous markups in the theoretical trade literature.²⁷

Demand Shocks and Product Reallocations

We have already described how an increase in demand L^c induces an increase in the toughness of competition λ in both the long-run (GE and PE) and the short-run. We now highlight how the demand conditions (B1) and (B1') generate a link between increases in the toughness of competition and product reallocations towards better performing products.

First, we note that the increase in competition λ induces a downward shift in output sales $q(v, \lambda)$ per-consumer (though not necessarily overall as the number of consumers is increasing). This decrease in output sales $q(v, \lambda)$ in turn generates changes in the price and marginal revenue elasticities ε_p and ε_ϕ , which depend on conditions (B1) and (B1'). Changes in those two elasticities then determine the changes in the elasticities of output, sales, and profit with respect to marginal cost v (see 8-10), which govern the reallocation of output, sales, and profit across firms with different costs v . We can now determine exactly how conditions (B1) and (B1') affect these reallocations:

Proposition 1 *(B1) is a necessary and sufficient condition for a positive demand shock to reallocate operating profits to better performing products: $\pi(v_1, \lambda)/\pi(v_2, \lambda)$ increases whenever $v_1 < v_2$.*

Proof. $|\varepsilon_{\pi, v}|$ increases for all products v whenever λ increases if and only if $\varepsilon_p(q_i)$ is increasing (see 10) ■

Proposition 2 *(B1') is a necessary and sufficient condition for a positive demand shock to reallocate output to better performing products: $q(v_1, \lambda)/q(v_2, \lambda)$ increases whenever $v_1 < v_2$.*

²⁶We will also show how the necessary condition for the French exporters evidence on product reallocations is slightly less restrictive than (B1').

²⁷Those functional forms include quadratic (linear demand), Bulow-Pfleiderer, CARA, and bipower preferences. In a separate Appendix, we review those functional forms and describe the parameter restrictions associated with conditions (B1) and (B1').

Proof. $|\varepsilon_{q,v}|$ increases for all products v whenever λ increases if and only if $\varepsilon_\phi(q_i)$ is increasing (see 8) ■

Proposition 3 (B1') is a sufficient condition for a positive demand shock to reallocate revenue to better performing products: $r(v_1, \lambda)/r(v_2, \lambda)$ increases whenever $v_1 < v_2$. The necessary condition is that $[1 - \varepsilon_p(q_i)]/\varepsilon_\phi(q_i)$ is decreasing.

Proof. $|\varepsilon_{r,v}|$ increases for all products v whenever λ increases if and only if $[1 - \varepsilon_p(q_i)]/\varepsilon_\phi(q_i)$ is decreasing (see 10). (B1') implies that $[1 - \varepsilon_p(q_i)]/\varepsilon_\phi(q_i)$ is decreasing.²⁸ ■

We have derived these reallocations using the per-consumer measures of performance, but since they are all evaluated as ratios, multiplying those by the number of consumers L^c would lead to identical outcomes (even though L^c is changing). Thus, we see that (B1') is a sufficient condition for all performance measures (profit, output, revenue) to be reallocated toward better performing products. In this case, an increase in competition (higher λ) induces a steeper relationship between a product's cost v and its profit, output, and revenue outcome (higher elasticities $|\varepsilon_{\pi,v}|, |\varepsilon_{q,v}|, |\varepsilon_{r,v}|$). A given percentage reduction in cost v then translates into a higher percentage increase in those performance outcomes. In the case of C.E.S. preferences, all those performance elasticities would be constant, and hence changes in demand (and corresponding changes in competition λ) would have no effect on the relative performance of products (conditional on selection into production).

The reallocation of output towards better performing products has a direct consequence for firm productivity: the allocation of firm employment to products must respond proportionately to the product-level output changes. Thus, for a given set of products, average productivity (an employment weighted average of product productivity $1/v$) must increase whenever output is reallocated towards better performing products.

Selection

In the short-run, we have already discussed how an increase in demand L^c and the corresponding increase in competition λ induces an increase in the cutoffs $\hat{c} = \hat{v}$. In this scenario, net profit per-product $L^c\pi(v, \lambda) - f$ is increasing for the high cost products with cost v close to the cutoff, even though the operating profit *per-consumer* $\pi(v, \lambda)$ is decreasing for all products (the increase in demand L^c dominates the negative impact of the increase in competition λ). Condition (B1)

²⁸Bertoletti and Epifani (2014) derive a similar expression for the reallocation of revenue. However, they do not show how that condition is implied by assumption (B1') and requires (B1) to hold.

then implies that the profits for the better performing products (with lower cost v) increase disproportionately relative to the high cost products. Thus, net operating profit per-product increases for all products.

In the long-run, we mentioned that the change in the cutoffs in response to an increase in demand could not be determined without making additional assumptions on demand. Under (B1), such a demand increase induces a disproportionate increase in operating profits for the best performing products; thus, total firm profit $\Pi(c, \lambda)$ becomes steeper (as a function of firm competency c). The free entry condition (12) then requires a single crossing of the new steeper total firm profit $\Pi(c, \lambda)$ curve with the old flatter one (ensuring that average profit in both cases is still equal to the constant entry cost f^e). This crossing defines a new profitability cutoff \tilde{c} whereby better performing firms with $c < \tilde{c}$ enjoy a profit increase whereas worse performing firms with $c > \tilde{c}$ suffer a profit loss. Hence, the zero profit cutoff \hat{c} (and the associated product-level cutoff \hat{v}) decreases, leading to the exit of the worse performing firms (and the worse performing products for all firms). The product-level net-profit curve $L^c\pi(v, \lambda) - f$ (as a function of product cost v) rotates in a similar fashion to the firm-level profit curve: profits for the best performing products increase while they decrease for the worse performing products. For those high performing products with low cost v , the increase in demand L^c dominates the effect of tougher competition (higher λ) on per-consumer profits $\pi(v, \lambda)$; whereas the opposite holds for the low performing products.

We further note that (B1) is also a necessary condition for this selection effect in the long run: If (B1) were violated then the profit curves would rotate in the opposite direction, reducing (increasing) profits for the best (worst) performing firms and products. This would result in an increase in the cutoffs $\hat{c} = \hat{v}$. In the limiting C.E.S. case, the cutoffs would be unaffected by changes in demands: the increase in demand L^c is exactly offset by the increase in competition leaving net-profits $L^c\pi(v, \lambda) - f$ unchanged for all products.

4.2 Open Economy

With our empirical findings in mind, we consider a simplified three-country economy consisting of a Home country (H : France) and a Foreign country (F : Rest of World) both exporting to a Destination country (D). We focus on the equilibrium in this destination D , characterized by a demand level L_D^c . For simplicity, we assume that this destination market is ‘small’ with respect to economy-wide outcomes in H and F , in the sense that changes in destination D -specific variables do

not affect those equilibrium outcomes (apart from those related to exports to D).²⁹ In particular, this entails that the number of entering firms in H and F , N_H^e and N_F^e , do not respond to changes in demand L_D^c in D . (Naturally, the number of entrants N_D^e into D will respond to demand conditions L_D^c in the long-run.)

Since we are interested in highlighting the impact of demand changes in destination D for competition there, we make one further simplification: We assume that firms in D do not export to either H or F . In this case, trade is necessarily unbalanced so we restrict our analysis to the PE case and assume an exogenous unitary wage. In the appendix, we show how to extend the equilibrium conditions to the GE and PE cases where firms in D export; and highlight how this would not change any of the qualitative results regarding the impact of demand shocks in D .

Exports from $l \in \{H, F\}$ into D incur both a per-unit iceberg cost $\tau_{lD} > 1$ as well as a fixed export cost f_{lD} . In order to streamline the analysis of the production decisions for domestic producers selling in D along with exporters into D , we use the notation $f_{DD} < f_{lD}$ to denote the smaller fixed cost faced by domestic producers and introduce $\tau_{DD} \equiv 1$. Lastly, we can allow for arbitrary technology differences across countries $l \in \{H, F, D\}$ in terms of the product ladder cost $z_l(m)$, the entry cost f_l^e , and the firm distribution of core competencies $\Gamma_l(c)$.

Product-Level Performance and Selection

We now consider the production decisions for sales into D . Since the marginal cost of production is constant and there are no cross-market or cross-product cost synergies, firms will make independent market-level decisions per-product. Any product i sold in D will face the residual inverse demand (4) with a unique endogenous competition level λ_D for that market. Any firm from $l \in \{H, F, D\}$ owning the blueprint for product i with marginal cost v faces a constant delivered marginal cost $\tau_{lD}v$ for market D . All these firms therefore solve the same profit maximization problem with respect to this delivered cost, associated with the same first order condition (6) as in the closed economy. Thus, we can use the same optimal output, operating profit, and revenue functions as in the closed economy to evaluate the performance of a product in destination D : $q(\tau_{lD}v, \lambda_D)$, $\pi(\tau_{lD}v, \lambda_D)$, $r(\tau_{lD}v, \lambda_D)$ capture those performance metrics (per-consumer in D) for a product with delivered cost $\tau_{lD}v$. Producers from $l \in \{H, F, D\}$ will sell a product with cost v in D so long as the associated profits are non-negative. This leads to three (for each origin) new zero-profit

²⁹This small open economy setup with monopolistic competition follows closely the development in Demidova and Rodrigues-Clare (2013) for the case of C.E.S. preferences.

cutoff conditions:

$$\pi(\tau_{lD}\hat{v}_{lD}, \lambda_D) = f_{lD}, \quad l \in \{H, F, D\}. \quad (14)$$

All products from l with cost $v \leq \hat{v}_{lD}$ will be sold in D . Since we have assumed that the fixed export cost is higher than the domestic market fixed cost ($f_{lD} > f_{DD}$ for $l \in \{H, F\}$), the marginal exported product will have a lower delivered cost than the marginal domestically produced product: $\tau_{lD}\hat{v}_{lD} < \hat{v}_{DD}$; and it will therefore feature strictly higher performance measures (in terms of operating profit, output, and revenues).

These product-level cutoffs will also be equal to the firm-level cutoffs for the sale of any product in D : Only firms with core competency $c \leq \hat{c}_{lD} = \hat{v}_{lD}$ will operate (sell) in market D . When $l \in \{H, F\}$, those cutoffs represent the firm-level export cutoff into D . When $l = D$, the cutoff represents the production (survival) cutoff for domestic firms in D .

Free Entry in the Long Run

The free entry conditions for markets H and F are irrelevant for destination D , due to our assumption of the former being ‘small’ relative to the latter two. In the long run, free entry into D equalizes the expected profits for prospective entrants with the sunk entry cost. Thus, this new free entry condition is identical to the one for the closed economy; except for the labeling of country- D specific variables:

$$\sum_{m=0}^{\infty} \left[\int_0^{\hat{c}_{DD}/z_D(m)} [\pi(cz_D(m), \lambda_D) L_D^c - f_{DD}] d\Gamma_D(c) \right] = f_D^e. \quad (15)$$

As in the closed economy case, this free entry condition, along with the zero cutoff profit condition (14, for $l = D$) jointly determine the equilibrium cutoffs $\hat{v}_{DD} = \hat{c}_{DD}$ along with the level of competition λ_D . The export cutoffs \hat{v}_{lD} for $l \in \{H, F\}$ are then determined by their matching zero cutoff profit condition in (14) given the competition level λ_D .

In the open economy, consumers in D spread their income over all the domestically produced and imported products. Since the number of entrants into H and F are exogenous with respect to country D , the budget constraint for consumers in D can still be used to directly determine the endogenous number of entrants into D (as was the case for the closed economy):

$$\sum_{l=H,F,D} \left(N_l^e \left\{ \sum_{m=0}^{\infty} \left[\int_0^{\hat{c}_{lD}/z_l(m)} r(\tau_{lD}cz_l(m), \lambda_D) d\Gamma_l(c) \right] \right\} \right) = 1. \quad (16)$$

Short-Run Equilibrium

In the short-run, the number of incumbents in D is fixed at \bar{N}_D (with the same exogenous distribution of core competencies $\Gamma_D(c)$); and the free entry condition for firms in D no longer holds. The budget constraint (16) still holds with the exogenous number of incumbents \bar{N}_D replacing the endogenous number of entrants N_D^e . Together with the three cutoff-profit conditions (14), these four equilibrium conditions determine the competition level λ_D along with the three cutoffs for $l \in \{H, F, D\}$.

Demand Shocks and Product Reallocations

Using the same reasoning as previously developed for the closed economy case, we can show that an increase in demand L_D^c will result in an increase in competition λ_D in D for both the long-run and short-run.³⁰ Thus, all of our previous results regarding the impact of such a demand shock on the reallocation of profits, output, and revenue towards better performing products still hold (Propositions 1-3). Thus, for exporters to D , we can connect demand conditions (B1) and (B1') to the reallocation of export sales and profits (in market D) towards better performing products.

Selection

Again using a similar reasoning to our analysis of the closed economy case, demand condition (B1) is sufficient for the demand shock to increase the net profit for the sales of the best performing products in D (the profits generated by sales in D) in both the long-run and short-run. Thus, so long as the fixed export cost f_{HD} is high enough, all exported products from H will fall into this category and experience a profit increase following the increase in demand. This, in turn, implies that an increase in demand leads to a fall in the export cutoffs \hat{v}_{HD} : Existing exporters from H increase their range of exported products to D , and some firms from H start exporting to D .

In the long-run, condition (B1) is also necessary for the demand shock to induce the selection of newly exported products from H to D (so long as the fixed export cost f_{HD} is high enough). Similar to the closed economy case, a violation of condition (B1) would imply a reduction in the profits of the best performing products to D , thus reversing the predicted consequences for export market selection.

³⁰Once again, a (weak) decrease in competition λ_D would necessarily lead to a violation of the free entry condition in the long-run, or the budget constraint in the short-run.

4.3 Connecting Back to Empirical Measures of Product Reallocations in Export Markets

We have just shown how our open economy model with demand condition (B1') can explain all of the evidence on the response of French exporters to demand shocks that we documented in Section 3. It explains how positive demand shocks induce the entry of new exported products and the reallocation of output and revenues towards the best performing products. The reallocation of output contributes positively to a firm's productivity (by shifting employment shares towards products with higher marginal products).³¹ And the reallocation of revenues generates an increase in the skewness of a firm's export sales to that destination.³² Demand condition (B1') and the weaker version (B1) are also directly connected to the empirical evidence on firm/product markups and pass-through. In the limiting case of C.E.S. preferences, demand shocks in export market would have no impact on the skewness of export sales; markups are constant across products; and pass-through is complete (equal to 1) for all products.

On its own, the evidence on the positive relationship between demand shocks and export skewness requires that $[1 - \varepsilon_p(q_i)] / \varepsilon_\phi(q_i)$ is decreasing over the range of exported output q_i that we observe. We have pointed out that condition (B1') ($\varepsilon_\phi(q_i)$ increasing) is sufficient for this outcome. However, since it is not a necessary condition, our evidence for the skewness of exports does not imply that (B1') must hold. On the other hand, we show in the appendix that the weaker condition (B1) (Marshall's Second Law of Demand, $\varepsilon_\phi(q_i)$ increasing) must nevertheless hold. In particular, we show that even if (B1) were violated over a portion of the relevant demand curve ($\varepsilon_p(q_i)$ decreasing over some range), then this would result in a reverse prediction for export skewness over this portion of the demand curve. This result also fits with what we know about the limiting case of C.E.S. preferences, where demand shocks have no impact on the skewness of export sales.

Thus, our empirical evidence on the impact of demand shocks for export skewness provides an independent confirmation for the empirical relevance of this critical property of demand – without relying on the measurement of prices and markups.

³¹We document the strong empirical connection between demand shocks and firm productivity in Section 6.

³²We showed how the ratio of export sales for any two products (with the better performing product in the numerator) increases in response to a demand shock. This clearly increases the skewness of export sales. In a separate appendix, we confirm that such an increase in skewness is reflected in the Theil and Atkinson indices that we use in our empirical work.

5 Trade Competition and Product Reallocations at the Firm-Level

Our theoretical model highlights how our measured demand shocks induce increases in competition for exporters to those destinations; and how the increased competition generates increases in productivity by shifting market shares and employment towards better performing products. We seek to directly measure this connection between demand shocks and productivity. Since we cannot measure the productivity associated with products sold to a particular destination, we need to show that the connection between demand shocks and product reallocations aggregates to the firm-level – before examining the link with firm-level productivity changes (which we can directly measure). Our results in section 3 highlighted how demand shocks lead to reallocations towards better performing products at the destination-industry level. In this section, we show how the destination-industry demand shocks can be aggregated to the firm-level – and this firm-level demand shock strongly predicts product reallocations towards better performing products (higher market shares) at the *firm-level*; that is, changes in skewness to the firm’s global product mix (the distribution of product sales across all destinations).

Intuitively, since there is a stable ranking of products at the firm level (better performing products in one market are most likely to be the better performing products in other markets – as we previously discussed), then reallocations towards better performing products within destinations should also be reflected in the reallocations of global sales/production towards better performing products; and the strength of this link between the skewness of sales at the destination and global levels should depend on the importance of the destination in the firm’s global sales. Our chosen measure of skewness, the Theil index, makes this intuition precise. It is the only measure of skewness that exhibits a stable decomposition from the skewness of global sales into the skewness of destination-level sales (see Jost 2007).³³ Specifically, let $T_{i,t}$ be firm i ’s Theil index for the skewness of its global exports by product $x_{i,t}^s \equiv \sum_d x_{i,d,t}^s$ (the sum of exports for that product across all destinations).³⁴ Then this global Theil can be decomposed into a market-share weighted average of the within-destination Theils $T_{i,d,t}$ and a “between-destination” Theil index $T_{i,d,t}^B$ that

³³This decomposition property is similar – but not identical – to the within/between decomposition of Theil indices across populations. In the latter, the sample is split into subsamples. In our case, the same observation (in this case, product sales) is split into “destinations” and the global measure reflects the sum across “destinations”.

³⁴The Theil index is defined in the same way as the destination level Theil in (2). We return to our notation $x_{i,d,t}^s$ to denote firm i ’s exports of product s to destination d in year t , which we use throughout our empirical analysis. In terms of our theoretical model, this corresponds to the export revenues generated for a product m by a firm with core competency c to destination D .

measures differences in the distribution of product-level market shares across destinations:³⁵

$$T_{i,t} = \sum_d \frac{x_{i,d,t}}{x_{i,t}} T_{i,d,t} - \sum_d \frac{x_{i,d,t}}{x_{i,t}} T_{i,d,t}^B, \quad (17)$$

where $x_{i,d,t} \equiv \sum_s x_{i,d,t}^s$ and $x_{i,t} \equiv \sum_d x_{i,d,t}$ represent firm i 's total exports to d , and across destinations to the world (global exports). The between-destination Theil $T_{i,d,t}^B$ is defined as

$$T_{i,d,t}^B = \sum_s \frac{x_{i,d,t}^s}{x_{i,d,t}} \log \left(\frac{x_{i,d,t}^s/x_{i,d,t}}{x_{i,t}^s/x_{i,t}} \right).$$

Note that the weights used in this decomposition for both the within- and between-destination Theils are the firm's export shares $x_{i,d,t}/x_{i,t}$ across destinations d . The between-destination Theil $T_{i,d,t}^B$ measures the deviation in a product's market share in a destination d , $x_{i,d,t}^s/x_{i,d,t}$, from that product's global market share $x_{i,t}^s/x_{i,t}$ and then averages these deviations across destinations. It is positive and converges to zero as the distributions of product market shares in different destination become increasingly similar.

To better understand the logic behind (17), note that it implies that the average of the within-destination Theil indices can be decomposed into the sum of two positive elements: the global Theil index, and the between-destination Theil index. This decomposition can be interpreted as a decomposition of variance/dispersion. The dispersion observed in the destination level product exports must be explained either by dispersion in global product exports (global Theil index), or by the fact that the distribution of product sales varies across destinations (between-destination Theil index).

A simple example helps to clarify this point. Take a firm with 2 products and 2 destinations. In each destination, exports of one product are x , and exports of the other product are $2x$. This leads to the same value for the within-destination Theil indices of $(1/3)\ln(1/3) + (2/3)\ln(2/3)$, and hence the same value for the average within-destination Theil index. Hence, if the same product is the better performing product in each market (with $2x$ exports), then the distributions will be synchronized across destinations and the between-destination Theil will be zero: all of the dispersion is explained by the global Theil index, whose value is equal to the common value of the two within-destination Theil indices. On the other hand, if the opposite products perform better in each market, global sales are $3x$ for each product. There is thus no variation in global product

³⁵For simplicity, we omit the industry referencing I for the destination Theils. The decomposition across industries follows a similar pattern.

sales, and the global Theil index is zero. Accordingly, all of the variation in the within-destination Theil indices is explained by the between-destination Theil.

The theoretical model of Bernard, Redding and Schott (2011) with CES demand predicts that the between-destination Theil index would be exactly zero when measured on a common set of exported products across destinations. With linear demand, Mayer, Melitz and Ottaviano (2014) show (theoretically and empirically) that this between-destination index would deviate from zero because skewness varies across destinations. We have shown earlier that this result holds for a larger class of demand systems such that the elasticity and the convexity of inverse demand increase with consumption. Yet, even in these cases, the between-destination Theil is predicted to be small because the ranking of the product sales is very stable across destinations. This leads to a prediction that the market-share weighted average of the destination Theils should be strongly correlated with the firm’s global Theil. Empirically, this prediction is strongly confirmed as shown in the firm-level scatter plot (across all years) in Figure 2.

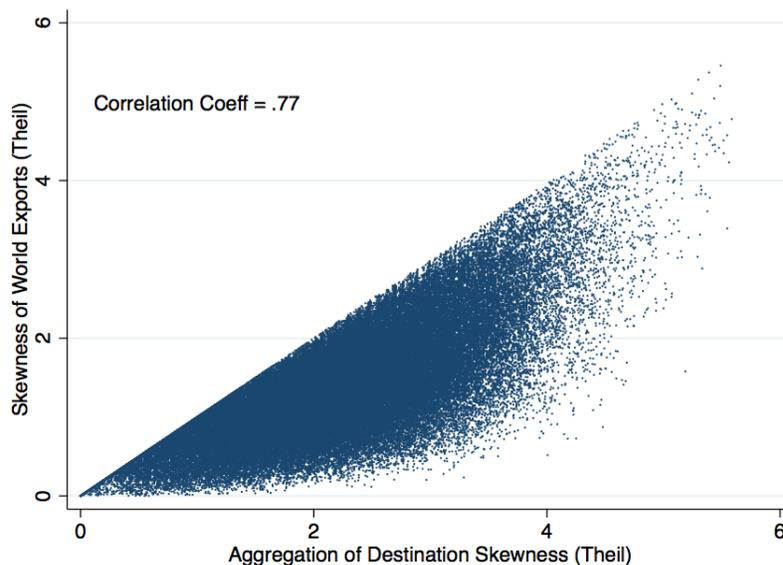


Figure 2: Correlation Between Global Skewness and Average Local Skewness

This high correlation between destination and global skewness of product sales enables us to move from our previous predictions for the effects of the demand shocks on skewness at the destination-level to a new prediction at the firm-level. To do this, we aggregate our destination-industry measures of demand shocks to the firm-level using the same weighing scheme by the firms’ export shares across destinations. We thus obtain our firm-level demand shock in (log) levels and

first difference:

$$\text{shock}_{i,t} \equiv \sum_{d,I} \frac{x_{i,d,t_0}^I}{x_{i,t_0}} \times \text{shock}_{i,d,t}^I, \quad \tilde{\Delta}\text{shock}_{i,t} = \sum_{d,I} \frac{x_{i,d,t-1}^I}{x_{i,t-1}} \times \tilde{\Delta}\text{shock}_{i,d,t}^I,$$

where $x_{i,t} \equiv \sum_{d,I} x_{i,d,t}^I$ represents firm i 's total exports in year t . As was the case for the construction of our firm-level destination shock (see equation 1), we only use the firm-level information on exported products and market shares in prior years (the year of first export sales t_0 for the demand shock in levels and lagged year $t - 1$ for the first difference between t and $t - 1$). This ensures the exogeneity of our constructed firm-level demand shocks (exogenous to firm-level actions in year $t > t_0$ for levels, and exogenous to firm-level changes Δ_t for first differences). In particular, changes in the set of exported products or exported market shares are not reflected in the demand shock.³⁶

By construction, our firm-level demand shocks will predict changes in the weighted average of destination skewness $T_{i,d,t}$ – and hence will predict changes in the firm's global skewness $T_{i,t}$ (given the high correlation between the two indices). This result is confirmed by our regression of the firms' global Theil on our trade shock measures, reported in the first three columns of Table 6. Our firm-level trade shock has a strong and highly significant (again, well beyond the 1% significance level) impact on the skewness of global exports. The industry level trade shock – which was already substantially weaker than the firm-level trade shock in the destination-level regressions – is no longer significant at the firm level. The GDP shock is not significant in the (log) levels regressions, but is very strong and significant in the two first-difference specifications.

Our global Theil measure $T_{i,t}$ measures the skewness of export sales across all destinations, but it does not entirely reflect the skewness of production levels across the firm's product range. That is because we cannot measure the breakdown of product-level sales on the French domestic market. Ultimately, it is the distribution of labor allocation across products (and the induced distribution of production levels) that determines a firm's labor productivity – conditional on its technology (the production functions for each individual product). As highlighted by our theoretical model, the export market demand shocks generate two different types of reallocations that both contribute to an increased skewness of production levels for the firm: reallocations within the set of exported products, which generate the increased skewness of global exports that we just discussed; but also reallocations from non-exported products towards the better performing exported products

³⁶Lileeva and Trefler (2010), Hummels et al (2014), and Bernard et al (2014) use a similar strategy to construct firm-level trade-related trade shocks.

Table 6: The Impact of Demand Shocks on the Global Product Mix (Firm Level)

Dependent Variable Specification	$T_{i,t}$		$\Delta T_{i,t}$		Exp. Intens $_{i,t}$	
	FE	FD	FD-FE	FE	FD	FD-FE
log GDP Shock	-0.001 (0.004)			0.003 ^a (0.001)		
log trade shock	0.045 ^a (0.009)			0.010 ^a (0.003)		
log trade shock - ISIC	-0.001 (0.001)			0.000 (0.000)		
$\tilde{\Delta}$ GDP Shock		0.117 ^a (0.031)	0.106 ^a (0.038)		0.041 ^a (0.010)	0.038 ^a (0.013)
$\tilde{\Delta}$ trade shock		0.057 ^a (0.011)	0.050 ^a (0.013)		0.016 ^a (0.003)	0.014 ^a (0.004)
$\tilde{\Delta}$ trade shock - ISIC		-0.003 (0.005)	-0.009 (0.007)		0.002 (0.002)	0.001 (0.002)
Observations	117981	117981	117981	111860	109049	109049

FE refers to firm-level fixed effects. Standard errors (clustered at the firm level) in parentheses: ^c < 0.1, ^b < 0.05, ^a < 0.01.

(including the extensive margin of newly exported products that we documented at the destination-level). Although we cannot measure the domestic product-level sales, we can measure a single statistic that reflects this reallocation from non-exported to exported goods: the firm's export intensity. We can thus test whether the demand shocks also induce an increase in the firm's export intensity. Those regressions are reported in the last three columns of Table 6, and confirm that our firm-level trade shock has a very strong and highly significant positive impact on a firm's export intensity.³⁷ The impact of the GDP coefficient is also strong and significant, whereas the industry-level trade shock remains insignificant. Thus, our firm-level trade shock and GDP shock both predict the two types of reallocations towards better performing products that we highlighted in our theoretical model (as a response to increased competition in export markets).

³⁷Since the export intensity is a ratio, we do not apply a log-transformation to that variable. However, specifications using the log of export intensity yield very similar results.

6 Trade Competition and Productivity

We just showed that our firm-level measure of demand shocks (aggregated across destinations) predicts increases in the skewness of global exports, and increases in export intensity. Holding firm technology fixed (the productivity of each individual product), this increase in the skewness of global production will generate productivity increases for the firm. We now directly test for this connection from the demand shocks to firm productivity.

We obtain our measure of firm productivity by merging our firm-level trade data with firm-level production data. This latter dataset contains various measures of firm outputs and inputs. As we are interested in picking up productivity fluctuations at a yearly frequency, we focus on labor productivity. We then separately control for the impact of changes in factor intensities and returns to scale (or variable utilization of labor) on labor productivity.

We compute labor productivity at the firm-level as deflated value added per worker assuming a sector-specific price deflator P^I . Note that this measure aggregates to the overall deflated value-added per worker for manufacturing. This aggregate productivity measure accurately tracks a welfare-relevant quantity index – even though we do not have access to firm-level prices: The effect of pure markup changes at the industry level are netted-out of our productivity measure.³⁸

More formally, we can write the welfare-relevant aggregate industry labor productivity – the ratio of industry deflated value-added (VA^I/P^I) over industry employment L^I – as the labor share weighted average of firm productivity using that same industry price deflator P^I for firm revenue:

$$\Phi^I \equiv \frac{VA^I/P^I}{L^I} = \sum_{i \in I} \frac{L_i}{L^I} \frac{VA_i/P^I}{L_i} \quad (18)$$

In other words, the revenue-based firm productivity measure $(VA_i/P^I)/L_i$ aggregates up to a quantity-based industry measure, using the empirically observed firm labor shares L_i/L^I and without any need for a quantity based output or productivity measure at the level of the firm (which would require measures of firm-product-level prices and qualities). This implicitly assumes the existence of an industry price aggregator P^I , though its functional form is left completely unrestricted. Consequently, we run all of our specifications with sector-time (2 digit NACE) fixed effects, thus eliminating the need for any direct measures of those sector-level deflators.³⁹ Our productivity

³⁸At the firm-level, an increase in markups across all products will be picked-up in our firm productivity measure – even though this does not reflect a welfare-relevant increase in output. But if this is the case, then this firm’s labor share will decrease, and its productivity will carry a smaller weight in the aggregate index.

³⁹We have also experimented with an alternative procedure using deflated value-added per worker (the value added

results therefore capture within-sector effects of the demand shocks, over-and-above any contribution of the sector deflator to a common productivity change across firms. We will thus report a welfare-relevant aggregate productivity change by aggregating our firm-level productivity changes using the observed changes in labor shares.

Our firm-level demand shocks only aggregate across export destinations. They therefore do not incorporate a firm’s exposure to demand shocks in its domestic (French) market. This is not possible for two reasons: most importantly, we do not observe the product-level breakdown of the firms’ sales in the French market (we only observe total domestic sales across products); in addition, world exports into France would not be exogenous to firm-level technology changes in France. Therefore, we need to adjust our export-specific demand shock using the firm’s export intensity to obtain an overall firm-level demand shock relevant for overall production and hence productivity:

$$\text{shock_intens}_{i,t} = \frac{x_{i,t_0}}{x_{i,t_0} + x_{i,F,t_0}} \text{shock}_{i,t}, \quad \tilde{\Delta} \text{shock_intens}_{i,t} = \frac{x_{i,t-1}}{x_{i,t-1} + x_{i,F,t-1}} \tilde{\Delta} \text{shock}_{i,t},$$

where $x_{i,F,t}$ denotes firm i ’s total (across products) sales to the French domestic market in year t (and the ratio thus measures firm i ’s export intensity).⁴⁰ Once again, we only use prior year’s information on firm-level sales to construct this overall demand shock. Note that this adjustment using export intensity is equivalent to assuming a demand shock of zero in the French market and including that market in our aggregation by market share relative to total firm sales $x_{i,t} + x_{i,F,t}$.

6.1 Impact of the Trade Shock on Firm Productivity

In this section, we investigate the direct link between this firm-level demand shock and firm productivity. Here, we focus exclusively on our firm-specific demand shock from (1). Most importantly, this is our only demand shock that exhibits variation across firms within destinations.⁴¹ In addition, this is the demand shock that induced the strongest and most robust response in the skewness of the intensive margin of product exports. Our measure of productivity is the log of value-added per

deflator coming from EUKLEMS dataset for France) and year fixed-effects (dropping the industry-year combined fixed effects). The differences in the results are negligible.

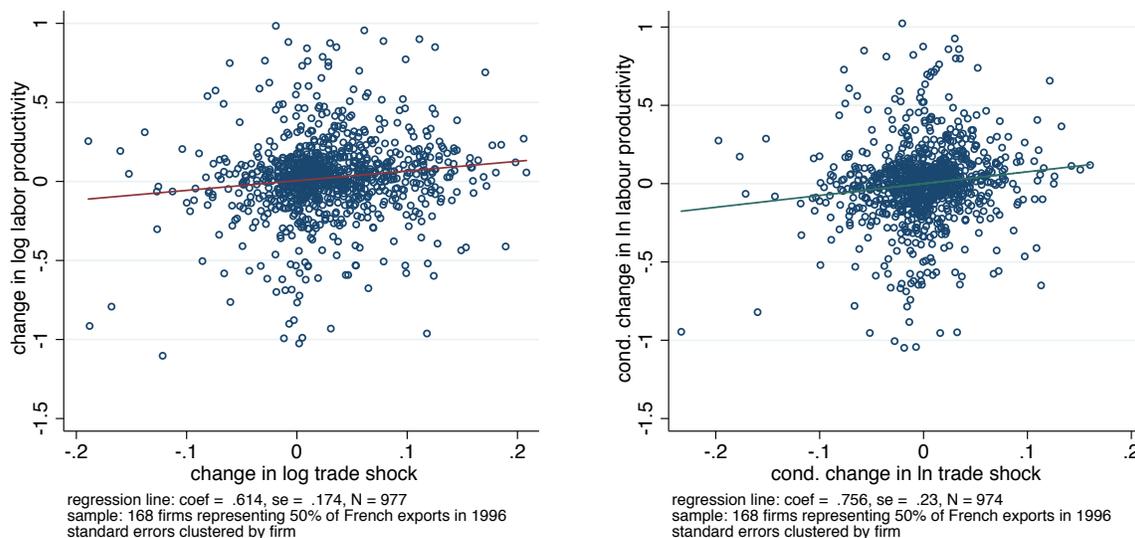
⁴⁰Since $x_{i,F,t}$ is only available in the firm balance sheet data (and not in the customs data), we use the reported firm total export figure $x_{i,t}$ from the same balance sheet data to compute this ratio. This ensures a consistent measurement for the firm’s export intensity. The correlation between this firm total export figure and the one reported by customs is .95 (very high, though exhibiting some differences between the two data sources).

⁴¹As a robustness check reported in appendix C, we purge the trade shock of any industry-destination-year effect in order to eliminate any correlation between destination characteristics and unobserved firm attributes associated with higher productivity growth. This approach precludes using the GDP and ISIC level trade shocks.

worker. All regressions include industry-year fixed effects, that will capture in particular different evolutions of price indexes across industries. In order to control for changes in capital intensity, we use the log of capital per worker. We also control for unobserved changes in labor utilization and returns to scale by using the log of raw materials (including energy use). Then, increases in worker effort or higher returns to scale will be reflected in the impact of raw materials use on labor productivity. As there is no issue with zeros for all these firm-level variables, we directly measure the growth rate of those variable using simple first differences of the log levels.

We begin with a graphical representation of the strong positive relationship between firm-level productivity and our constructed demand shock. Figure 3 illustrates the correlation between those variables in first differences for the largest French exporters (representing 50% of French exports in 1996). Panel (a) is the unconditional scatter plot for those variables, while panel (b) shows the added-variable plot for the first-difference regression of productivity on the trade shock, with additional controls for capital intensity, raw materials (both in log first-differences) and time dummies. Those figures clearly highlight the very strong positive response of the large exporters' productivity to changes in trade competition in export markets (captured by the demand shock).

Figure 3: Exporters Representing 50% of French Trade in 1996: First Differences 1996-2005



(a) Unconditional

(b) Conditional

Table 7 shows how this result generalizes to our full sample of firms and our three different specifications (FE, FD, FD-FE). Our theoretical model emphasizes how a multi-product firm's productivity responds to the demand shock via its effect on competition and product reallocations

in the firm’s export markets. Thus, we assumed that the firm’s technology at the product level (the marginal cost $v(m, c)$ for each product m) was exogenous (in particular, with respect to demand fluctuations in export markets). However, there is a substantial literature examining how this technology responds to export market conditions via various forms of innovation or investment choices made by the firm. We feel that the timing dimension of our first difference specifications – especially our FD-FE specifications which nets out any firm-level growth trends – eliminates this technology response channel: It is highly unlikely that a firm’s innovation or investment response to the trade shock in a given year (especially with respect to the trade shock’s deviation from trend growth in the FD-FE specification) would be reflected contemporaneously in the firm’s productivity. However, we will also show some additional robustness checks that address this potential technology response.

The first three columns of Table 7 show that, across our three timing specifications, there is a stable and very strong response of firm productivity to the trade shock. Since our measure of productivity as value added per worker incorporates neither the impact of changes in input intensities nor the effects of non-constant returns to scale, we directly control for these effects in the next set of regressions. In the last 3 columns of Table 7, we add controls for capital per worker and raw material use (including energy). Both of these controls are highly significant: not surprisingly, increases in capital intensity are reflected in labor productivity; and we find that increases in raw materials use are also associated with higher labor productivity. This would be the case if there are increasing returns to scale in the value-added production function, or if labor utilization/effort increases with scale (in the short-run). However, even when these controls are added, the very strong effect of the trade shock on firm productivity remains highly significant well beyond the 1% significance level (from here on out, we will keep those controls in all of our firm-level productivity regressions).

Although our trade shock variable exhibits variation within destinations it nevertheless also reflects the impact of aggregate destination shocks. This could potentially generate a selection bias if firms with unobserved characteristics leading to higher productivity growth self-select into destinations with positive aggregate destination shocks (or vice-versa). In order to eliminate this possibility, we generate a version of our trade shock that is purged (demeaned) of any industry-destination-year effects (trends for the FD-FE specification). Using this purged measure instead of our original one strengthen the impact of the trade shock on firm productivity (but leaves the impact of the other controls virtually un-changed). This robustness check is reported in the appendix.

Table 7: Baseline Results: Impact of Trade Shock on Firm Productivity

Dependent Variable Specification	log prod.			Δ log prod.		
	FE	FD	FD-FE	FE	FD	FD-FE
log (trade shock × export intens.)	0.061 ^a (0.016)			0.051 ^a (0.016)		
$\tilde{\Delta}$ (trade shock × export intens.)		0.106 ^a (0.019)	0.106 ^a (0.023)		0.112 ^a (0.020)	0.113 ^a (0.024)
log capital stock per worker				0.117 ^a (0.004)		
log raw materials				0.086 ^a (0.003)		
Δ log capital stock per worker					0.125 ^a (0.005)	0.133 ^a (0.006)
Δ log raw materials					0.092 ^a (0.003)	0.090 ^a (0.003)
Observations	213001	185688	185688	203977	175619	175619

FE refers to firm-level fixed effects. All regressions also include industry-year dummies. Standard errors (clustered at the firm level) in parentheses: ^c < 0.1, ^b < 0.05, ^a < 0.01.

We now describe several robustness checks that further single-out our theoretical mechanism operating through the demand-side product reallocations for multi-product firms. In the next table we regress our capital intensity measure on our trade shock; the results in Table 8 show that there is no response of investment to the trade shock. This represents another way to show that the short-run timing for the demand shocks precludes a contemporaneous technology response: if this were the case, we would expect to see some of this response reflected in higher investment (along with other responses along the technology dimension).

Next we use a different strategy to control for the effects of non-constant returns to scale or variable labor utilization: in Table 9, we split our sample between year intervals where firms increase/decrease employment. If the effects of the trade shock on productivity were driven by scale effects or higher labor utilization/effort, then we would expect to see the productivity responses concentrated in the split of the sample where firms are expanding employment (and also expanding more generally). Yet, Table 9 shows that this is not the case: the effect of the trade shock on productivity is just as strong (even a bit stronger) in the sub-sample of years where firms are decreasing employment; and in both cases, the coefficients have a similar magnitude to our baseline

Table 8: Capital Intensity Does Not Respond to Trade Shocks

Dependent Variable	$\ln K/L$	$\Delta \ln K/L$	$\Delta \ln K/L$
Specification	FE	FD	FD-FE
$\log(\text{trade shock} \times \text{export intens.})$	0.029 (0.021)		
$\tilde{\Delta}(\text{trade shock} \times \text{export intens.})$		0.015 (0.022)	0.002 (0.027)
Observations	218073	190512	190512

FE refers to firm fixed effects. All regressions also include industry-year dummies. Standard errors (clustered at the firm level) in parentheses: $c < 0.1$, $b < 0.05$, $a < 0.01$.

results in Table 7.⁴²

Table 9: Robustness to Scale Effects

Sample	Employment Increase	Employment Decrease
Dependent Variable	$\Delta \log$ productivity	$\Delta \log$ productivity
Specification	FD	FD
$\tilde{\Delta}(\text{trade shock} \times \text{export intens.})$	0.128 ^a (0.034)	0.170 ^a (0.033)
$\Delta \log$ capital stock per worker	0.108 ^a (0.006)	0.104 ^a (0.006)
$\Delta \log$ raw materials	0.100 ^a (0.004)	0.096 ^a (0.004)
Observations	69881	65739

All regressions include industry-year dummies. Standard errors (clustered at the firm level) in parentheses: $c < 0.1$, $b < 0.05$, $a < 0.01$.

A potential concern with our trade shock variable weighting demand shocks in the destination country by export shares is that it would be correlated with import shocks, i.e. with supply shocks originating from the same foreign countries and affecting directly production costs of French firms through imported intermediate goods. We therefore construct a symmetric set of variables that weight changes in exports (to the world except France) from a given product/country by the firm-level import shares from that product/country. Table 10 introduces those new variables into the baseline specification. Although import shocks have a separate large and very significant effect, it

⁴²Since we are splitting our sample across firms, we no longer rely on the two specifications with firm fixed-effects and only show results for the FD specification.

does not affect our main effect of interest in any major way.

Table 10: Robustness to Import Shocks

Dependent Variable Specification	log prod.		Δ log prod.		log prod.		Δ log prod.	
	FE	FD	FD-FE	FE	FD	FD-FE		
log (trade shock \times export intens.)	0.068 ^a (0.020)			0.056 ^a (0.019)				
log (trade shock \times import intens.)	0.078 ^a (0.029)			0.065 ^b (0.028)				
$\tilde{\Delta}$ (trade shock \times export intens.)		0.110 ^a (0.022)	0.111 ^a (0.026)		0.129 ^a (0.022)		0.133 ^a (0.027)	
$\tilde{\Delta}$ (trade shock \times import intens.)		0.225 ^a (0.034)	0.245 ^a (0.041)		0.197 ^a (0.033)		0.218 ^a (0.041)	
log capital stock per worker				0.095 ^a (0.004)				
log raw materials				0.090 ^a (0.004)				
Δ log capital stock per worker					0.107 ^a (0.005)		0.115 ^a (0.006)	
Δ log raw materials					0.091 ^a (0.003)		0.090 ^a (0.003)	
Observations	133382	152171	152171	129517	144938	144938		

FE refers to firm-level fixed effects. All regressions also include industry-year dummies. Standard errors (clustered at the firm level) in parentheses: ^c < 0.1, ^b < 0.05, ^a < 0.01.

In order to further single-out our theoretical mechanism operating through the demand-side product reallocations for multi-product firms, we now report two different types of falsification tests. Our first test highlights that the link between productivity and the trade shocks is only operative for multi-product firms. Table 11 reports the same regression (with controls) as our baseline results from Table 7, but only for single-product exporters. This new table clearly shows that there is no evidence of this link among this subset of firms. Next, we show that this productivity-trade link is only operative for firms with a substantial exposure to export markets (measured by export intensity). Similarly to single-product firms, we would not expect to find a significant productivity-trade link among firms with very low export intensity. This is indeed the case. In Table 12, we re-run our baseline specification using the trade shock before it is interacted

with export intensity. The first three columns report the results for the quartile of firms with the lowest export intensity, and highlight that there is no evidence of the productivity-trade link for those firms. On the other hand, we clearly see from the last three columns that this effect is very strong and powerful for the quartile of firms with the highest export intensity.⁴³

When we focus on this top quartile for export intensity, we are restricting our sample to substantially larger exporters. If a French exporter is dominant in some of its export destinations, then our exogeneity assumption for rest-of-world exports to those destinations may be violated. In this situation where a French firm has a substantial market share for some destination-product market, it is possible for the rest-of-the-world exports to respond to that firm's change in productivity. First, we note that this would likely bias our coefficient downwards, as a French firm's productivity increase would induce other exporters to reduce their shipments in their competing markets. We can nevertheless address this potential endogeneity directly by dropping any firm that could hypothetically impact rest-of-world exports. We have experimented with several different market thresholds, and have found that exclusions of those firms either leave those key coefficients on the trade shock unaffected, or sometimes raises them.⁴⁴

The firms with high export intensity therefore have a response of productivity to trade shocks estimated around 10% (columns 5 and 6 of Table 12). How should we interpret this number in terms of the impact of our mechanism on the productivity of the French economy as a whole? As we previously discussed, our firm-level measure of productivity does not correspond directly to physical productivity since we cannot adjust for changes in firm-product prices and qualities. However, as we pointed out (see equation 18), this firm-level productivity measure can be aggregated up to a sector level measure of physical productivity using the firms' labor shares. Side-stepping the issues of measurement for the sector price index P^I , we can directly measure the separate contribution of the trade shock to sectoral productivity.

We start with the coefficient on the trade shock for the high export-intensity quartile β_{q4} for the specifications in first differences. This gives us the percentage productivity response to a 1% trade shock for the firms in the fourth quartile. We assume that there is no response of productivity for the firms in the remaining 3 quartiles. We can then use each firm's measured trade shock to calculate an

⁴³Since the trade shock has not been interacted with export intensity, the coefficients for this top quartile represent significantly higher magnitudes than the average coefficients across the whole sample reported in Table 7 (since export intensity is always below 1). This is also confirmed by a specification with the interacted trade shock restricted to this same top quartile of firms.

⁴⁴For example, one of our more aggressive thresholds was the exclusion of any firm with an average product-destination market share (for any year) exceeding 10%. In this case, the coefficients for the trade shock corresponding to columns 4-6 of Table 12 are (standard errors in parentheses): FE: .070 (.030); FD: .170 (.038); FD-FE: .170 (.045).

Table 11: Robustness: Single Product Firms

Dependent Variable Specification	log prod.		Δ log prod.
	FE	FD	FD-FE
log (trade shock \times export intens.)	-0.050 (0.047)		
log capital stock per worker	0.180 ^a (0.013)		
log raw materials	0.085 ^a (0.007)		
$\tilde{\Delta}$ (trade shock \times export intens.)		0.038 (0.044)	0.012 (0.066)
Δ log capital stock per worker		0.214 ^a (0.013)	0.260 ^a (0.019)
Δ log raw materials		0.103 ^a (0.007)	0.100 ^a (0.011)
Observations	33198	25519	25519

FE refers to firm fixed effects. All regressions also include industry-year dummies. Standard errors (clustered at the firm level) in parentheses: ^c < 0.1, ^b < 0.05, ^a < 0.01.

induced percentage productivity change for each firm in the fourth quartile. Aggregating those using the observed changes in the firms' labor shares yields the overall contribution to aggregate physical productivity changes from the observed trade shocks (at the sector or total manufacturing level). Since equation (18) aggregates productivity in levels, we calculate the counterfactual productivity level for each firm in period $t + 1$ by applying this productivity change to the firm's observed productivity in period t . We then compute the ratios of the counterfactual aggregate productivity level for $t + 1$ and the observed aggregate productivity level in period t . The percentage changes reported in column 1 of Table 13 are the yearly averages of those ratios. This table uses the $\beta_{q4} = .092$ coefficient from the FD specification (the results for $\beta_{q4} = .101$ for the FD-FE specification would be roughly 10% higher). We use the same aggregation method to obtain sector and total manufacturing level changes in the trade shock – only for the firms in the fourth quartile (the trade shock for all remaining firms is set to 0). Those aggregate changes in the trade shock are reported in column 2. The labor shares used to compute those aggregate changes in columns 1 and 2 are based on total employment across all export intensity quartiles (not just as a fraction of employment in the fourth quartile). Those reported changes are therefore diluted by the zero contributions for all firms in the bottom three quartiles. Column 3 reports the employment share for those firms

Table 12: Robustness: Low/High Export Intensity

Sample Dependent Variable Specification	exp. intens. quartile # 1			exp. intens. quartile # 4		
	log prod. FE	Δ log prod. FD FD-FE		log prod. FE	Δ log prod. FD FD-FE	
log trade shock	0.003 (0.006)			0.072 ^a (0.011)		
log capital stock per worker	0.117 ^a (0.009)			0.104 ^a (0.007)		
log raw materials	0.070 ^a (0.005)			0.111 ^a (0.006)		
$\tilde{\Delta}$ trade shock		0.004 (0.008)	0.006 (0.009)		0.092 ^a (0.014)	0.101 ^a (0.016)
Δ log capital stock per worker		0.125 ^a (0.011)	0.129 ^a (0.013)		0.107 ^a (0.008)	0.114 ^a (0.010)
Δ log raw materials		0.084 ^a (0.006)	0.081 ^a (0.007)		0.108 ^a (0.006)	0.104 ^a (0.007)
Observations	38806	30909	30909	57267	48716	48716

Standard errors (clustered at the firm level) in parentheses: ^c < 0.1, ^b < 0.05, ^a < 0.01.

in the fourth quartile, while column 4 reports the overall employment share (all quartiles) for the sector (as with the first two columns those numbers are yearly averages). Our results for the productivity changes range from an average annual impact of 3.4 % for wearing apparel, to a small negative impact for refined petroleum for instance. Overall, we find that the aggregate impact of the trade shocks is substantial – accounting for a 1.2% average productivity growth rate for the entire French manufacturing sector (working only through the productivity linkages for the firms with the highest export intensities). This amounts to a 12% productivity gain over our ten year sample period from 1995-2005. Taking the ratio of the annual 1.2% productivity gain to the 6.2% trade shock change, we see that the .092 firm-level elasticity obtained in Table 12 is doubled to .193 for aggregate productivity. This reflects the contribution of labor share reallocations towards firms with relatively higher trade shocks.

7 Conclusion

This paper uses detailed firm-level data to assess the relevance and magnitude of a new channel of gains from trade: the productivity gains associated with demand shocks in export markets (via

Table 13: Quantification of Trade Shock Effect on Productivity

Industry	prod.	trade shock	% high exp.intens.	% mfg. emp.
Wearing apparel	3.38	5.21	27.36	2.26
Wood	3.37	6.34	20.36	1.70
Tobacco	3.22	43.60	.48	.16
Printing and publishing	2.81	8.48	5.36	3.31
Radio, television and communication	1.80	4.94	59.77	4.31
Leather and footwear	1.79	3.59	26.86	1.21
Textiles	1.69	1.99	33.04	3.29
Motor vehicles, trailers and semi-trailers	1.62	9.80	52.39	7.82
Machinery	1.32	5.54	45.40	9.12
Manufacturing nec	1.19	5.94	22.72	3.56
Pulp and paper	1.18	3.67	30.62	2.82
Chemicals	1.15	6.58	40.55	9.63
Fabricated metal	.94	7.04	17.41	8.81
Medical, precision and optical instruments	.85	5.84	46.82	3.53
Rubber and plastics	.80	5.75	36.97	7.18
Electrical machinery	.73	5.83	53.12	5.17
Basic metals	.70	6.27	58.91	4.06
Food and beverages	.66	6.20	14.12	11.88
Other transport equipment	.65	7.25	69.14	4.30
Office machinery	.64	3.70	42.55	1.09
Other non-metallic mineral	.46	3.89	35.52	3.86
Coke, refined petroleum and nuclear fuel	-.18	5.12	25.54	.93
Total Manufacturing	1.17	6.20	36.66	100.00

Columns (1) and (2) provide average percentage changes over the 1996-2005 period.

the induced effect of demand on the product mix of exporters). Our theoretical model predicts that demand shocks generate an endogenous increase in local competition that induces firms to skew their sales towards their better performing products – generating increases in productivity. Empirically, our data matches individual export flows by French firms to each country in the world with balance sheet data needed to evaluate the impact of demand shocks on productivity and to control for confounding factors. The strategy is therefore to look at product mix changes inside the firm, rather than reallocations of market shares across firms. We can therefore control for many alternative explanations that might be correlated with foreign demand shocks – a strategy that would not be possible when evaluating the effects *across* firms.

Our baseline results shows that the elasticity of labor productivity to trade shocks is between 5 and 11 %. This order of magnitude is very robust to controls for short-run investment by the firm, scale effects, and possibly correlated import shocks. Our measured productivity effect for

single product firms is nil, further highlighting the importance of changes in product mix for multi-product firms. We also show that this productivity response is concentrated within the quartile of exporters with the highest export intensities. Taking into account the weight of those firms in the whole economy, we calculate that the average annual increase in French manufacturing productivity – in response to growth in world trade – over our 10 year sample (from 1995-2005) is slightly over 1 percent per year.

References

- Amiti, Mary, Oleg Itskhoki, and Jozef Konings. 2014. “Importers, Exporters, and Exchange Rate Disconnect.” *American Economic Review* 104 (7): 1942–78.
- Arkolakis, Costas, Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare. 2015. “The Elusive Pro-Competitive Effects of Trade.” NBER Working Paper 21370.
- Asplund, Marcus, and Volker Nocke. 2006. “Firm Turnover in Imperfectly Competitive Markets.” *Review of Economic Studies* 73 (2): 295–327.
- Atkinson, Anthony B. and Andrea Brandolini. 2010. “On Analyzing the World Distribution of Income.” *World Bank Economic Review* 24 (1): 1-37.
- Baldwin, John R. and Wulong Gu. 2009. “The Impact of Trade on Plant Scale, Production-Run Length and Diversification.” In *Producer Dynamics: New Evidence from Micro Data*, edited by T. Dunne, J.B. Jensen, and M.J. Roberts. Chicago: University of Chicago Press.
- Berman, Nicolas, Philippe Martin, and Thierry Mayer. 2012. “How Do Different Exporters React to Exchange Rate Changes?” *Quarterly Journal of Economics* 127 (1): 437-492.
- Bernard, Andrew B., Andreas Moxnes, and Karen Helene Ulltveit-Moe. 2014. “Two-Sided Heterogeneity and Trade.” NBER Working Paper 20136.
- Bernard, Andrew B., Stephen J. Redding and Peter K. Schott. 2011. “Multi-product Firms and Trade Liberalization.” *Quarterly Journal of Economics* 126 (3): 1271-1318.
- Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott. 2012. “The Empirics of Firm Heterogeneity and International Trade.” *Annual Review of Economics* 4 (1): 283–313.
- Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding and Peter K. Schott. 2007. “Firms in International Trade.” *Journal of Economic Perspectives* 21 (3): 105-130.
- Bertoletti, Paolo, and Paolo Epifani. 2014. “Monopolistic Competition: CES Redux?” *Journal of International Economics* 93 (2): 227–38.
- Burstein, Ariel, and Gita Gopinath. 2015. “International Prices and Exchange Rates.” In *Handbook of International Economics*, edited by Kenneth Rogoff, Gita Gopinath, and Elhanan Helpman, Volume 4:391–451. Elsevier.

- Chatterjee, Arpita, Rafael Dix-Carneiro, and Jade Vichyanond. 2013. “Multi-Product Firms and Exchange Rate Fluctuations.” *American Economic Journal: Economic Policy* 5 (2): 77–110.
- Deaton, Angus, and John Muellbauer. 1980. *Economics and Consumer Behavior*. Cambridge University Press.
- De Loecker, Jan, Penny Goldberg, Amit Khandelwal and Nina Pavcnik. 2016. “Prices, Markups and Trade Reform.” *Econometrica* 84(2):445-510.
- Demidova, Svetlana and Andres Rodríguez-Clare. 2013. “The Simple Analytics of the Melitz Model in a Small Economy.” *Journal of International Economics*, 90 (2): 266–272.
- Dhingra, Swati, and John Morrow. 2014. “Monopolistic Competition and Optimum Product Diversity Under Firm Heterogeneity.” Unpublished, London School of Economics.
- Dhyne, Emmanuel, Amil Petrin, Valerie Smeets, and Frederic Warzynski. 2016. “Multi-Product Firms, Import Competition, and the Evolution of Firm-product Technical Efficiencies.” Unpublished, University of Aarhus.
- Fabinger, Michal, and E. Glen Weyl. 2014. “A Tractable Approach to Pass-Through Patterns with Applications to International Trade,” University of Chicago.
- Gaulier, Guillaume and Soledad Zignago. 2010. “BACI: International Trade Database at the Product-Level. The 1994-2007 Version.” CEPII Working Paper, N° 2010-23.
- Harrison, Ann E. 1994. “Productivity, Imperfect Competition and Trade Reform : Theory and Evidence.” *Journal of International Economics* 36 (1-2): 53-73.
- Hummels, David, Rasmus Jorgensen, Jakob Munch, and Chong Xiang. 2014. “The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data.” *American Economic Review* 104 (6): 1597–1629.
- Iacovone, Leonardo, and Beata S. Javorcik. 2010. “Multi-Product Exporters: Product Churning, Uncertainty and Export Discoveries.” *The Economic Journal* 120 (544): 481–99.
- Jost, Lou. 2007. “Partitioning Diversity into Independent Alpha and Beta Components.” *Ecology* 88 (10): 2427–39.
- Kimball, Miles S. 1995. “The Quantitative Analytics of the Basic Neomonetarist Model.” *Journal of Money, Credit, and Banking* 27 (4): 1241.
- Levinsohn, James A. 1993. “Testing the Imports-as-market-discipline Hypothesis.” *Journal of International Economics* 35 (1-2): 1-22.
- Li, Hongbin, Hong Ma, and Yuan Xu. 2015. “How Do Exchange Rate Movements Affect Chinese Exports? — A Firm-Level Investigation.” *Journal of International Economics* 97 (1): 148–61.
- Lileeva, Alla, and Daniel Trefler. 2010. “Improved Access to Foreign Markets Raises Plant-Level Productivity... For Some Plants.” *The Quarterly Journal of Economics* 125 (3): 1051–99.
- Mayer, Thierry and Gianmarco I.P. Ottaviano. 2008. “The Happy Few: The Internationalisation of European Firms.” *Intereconomics* 43 (3): 135-148.

- Mayer, Thierry, Marc J. Melitz, and Gianmarco I.P. Ottaviano. 2014. “Market Size, Competition, and the Product Mix of Exporters” , *American Economic Review* 104 (2): 495–536.
- Melitz, Marc J. and Gianmarco I.P. Ottaviano. 2008. “Market Size, Trade and Productivity.” *Review of Economic Studies* 75 (1): 295 -316.
- Melitz, Marc J., and Stephen J. Redding. 2014. “Heterogeneous Firms and Trade.” In Handbook of International Economics, edited by Elhanan Helpman, Kenneth Rogoff, and Gita Gopinath, 4:1–54. Elsevier.
- Melitz, Marc J., and Stephen J. Redding. 2015. “Heterogeneous Firms and Trade.” In , edited by Elhanan Helpman, Kenneth Rogoff, and Gita Gopinath, 4:1–54. Handbook of International Economics 4. Elsevier.
- Mrazova, Monika, and J. Peter Neary. 2014. “Not So Demanding: Preference Structure, Firm Behavior, and Welfare.” Unpublished, Oxford University.
- Parenti, Mathieu, Philip Ushchev and Jacques-François Thisse. 2014. “Toward a theory of monopolistic competition.” Unpublished, Université Catholique de Louvain.
- Zhelobodko, Evgeny, Sergey Kokovin, Mathieu Parenti and Jacques-François Thisse. 2012. “Monopolistic Competition: Beyond the Constant Elasticity of Substitution.” *Econometrica* 80 (6): 2765-2784.

Appendix

A Demand Conditions (B1) and (B1')

The definitions of inverse demand (4) and marginal revenue (5) imply that the elasticities of these two curves must satisfy:

$$\frac{\varepsilon'_p(q_i)q_i}{1 - \varepsilon_p(q_i)} = \varepsilon_\phi(q_i) - \varepsilon_p(q_i). \quad (19)$$

In addition, the assumptions (A1)-(A3) required for the consumer and firm maximization problems further imply that $0 \leq \varepsilon_p(q_i) < 1$ and $\varepsilon_\phi(q_i) > 0$. Thus, we see that $\varepsilon_p(q_i)$ is increasing (decreasing) if and only if $\varepsilon_\phi(q_i)$ is above (below) $\varepsilon_p(q_i)$. Thus, demand condition (B1) ($\varepsilon_p(q_i)$ increasing) holds if and only if the pass-through elasticity $\theta \equiv \partial \ln p(q(v, \lambda)) / \partial \ln v = \varepsilon_p / \varepsilon_\phi$ is less than 1 (pass-through is incomplete).

We now show that demand condition (B1') implies (B1) and that $[1 - \varepsilon_p(q_i)] / \varepsilon_\phi(q_i)$ decreasing also implies (B1). To show this, we impose one further – but very mild – condition on the convexity of demand at $q_i = 0$: $\lim_{q_i \rightarrow 0} \varepsilon'_p(q_i)q_i \geq 0$. This just rules out an extreme form of demand convexity at zero consumption (though, implicitly, this assumption also relies on the continuity of demand for $q_i \geq 0$). It can only be violated if $\lim_{q_i \rightarrow 0} \varepsilon'_p(q_i) = -\infty$ and in addition requires $\varepsilon'_p(q_i)$ to decrease to $-\infty$ faster than $-1/q_i$ as $q_i \rightarrow 0$. So long as this condition holds, then $\varepsilon_\phi(0) \geq \varepsilon_p(0)$ with equality if $\varepsilon'_p(0) \leq 0$.

A.1 (B1') Implies (B1)

Assume (B1') holds: $\varepsilon'_\phi(q_i) > 0$. Given (19), $\varepsilon'_p(q_i) \leq 0$ implies $\varepsilon_p(q_i) \geq \varepsilon_\phi(q_i)$ over that range for q_i . Given our assumption that $\varepsilon_p(0) \leq \varepsilon_\phi(0)$, this can only happen after $\varepsilon_p(q_i)$ cuts $\varepsilon_\phi(q_i)$ from below (or at $q_i = 0$). This in turn would imply that $\varepsilon_p(q_i)$ is strictly increasing at the intersection point $q_i = \bar{q}$: $\varepsilon_p(\bar{q}) \geq \varepsilon_\phi(\bar{q}) > 0$. This contradicts (19) which requires $\varepsilon'_p(\bar{q}) = 0$. Thus, $\varepsilon_p(q_i)$ must be strictly below $\varepsilon_\phi(q_i)$ for all $q_i > 0$. Given (19), this entails $\varepsilon'_p(q_i) \geq 0$ with the inequality strict for $q_i > 0$. Thus, $\varepsilon_p(q_i)$ must be strictly increasing.

A.2 $[1 - \varepsilon_p(q_i)] / \varepsilon_\phi(q_i)$ Decreasing Implies (B1)

If $[1 - \varepsilon_p(q_i)] / \varepsilon_\phi(q_i)$ is decreasing, then $\varepsilon_p(q_i)$ and $\varepsilon_\phi(q_i)$ cannot both be decreasing. However $\varepsilon'_p(q_i) < 0$ would imply that $\varepsilon_\phi(q_i)$ is also decreasing and thus represents a violation. More formally,

the average $\varepsilon'_\phi(q_i)$ over any interval $[0, \bar{q}]$ must then be negative (given (19) and $\varepsilon'_p(q_i) < 0$):

$$\overline{\varepsilon'_\phi} = \frac{1}{\bar{q}} \int_0^{\bar{q}} \varepsilon'_\phi(q_i) dq_i = \varepsilon_\phi(\bar{q}) - \varepsilon_\phi(0) \leq \varepsilon_p(\bar{q}) - \varepsilon_p(0) < 0.$$

This highlights that $\varepsilon'_\phi(q_i) \geq 0$ cannot hold over the entire interval $[0, \bar{q}]$. So $\varepsilon_\phi(q_i)$ is decreasing on average, and must be decreasing over a subset of $[0, \bar{q}]$ for any $\bar{q} > 0$. This reasoning also applies to cases where $\varepsilon_p(q_i)$ is initially increasing then decreasing. Say that $\varepsilon_p(q_i)$ is decreasing starting at $q_1 > 0$. Then, the average $\overline{\varepsilon'_\phi} < 0$ over any interval $[q_1, q_2]$. (Since the ε_ϕ and ε_p curves must cross at q_1 where $\varepsilon_p(q_1) = 0$.)

B Allowing for the Destination Country to Export

In the main text we simplified our analysis by assuming that firms in D do not export. In the short-run, allowing for such exporting opportunities would have no impact on the equilibrium outcome for D . In the long-run, the free entry condition for D would change, reflecting the average exporting profits for a prospective entrant:

$$\sum_{l=F,H,D} \left(\sum_{m=0}^{\infty} \left\{ \int_0^{\hat{c}_{Dl}/z_D(m)} [\pi(\tau_{Dl}cz(m), \lambda_l) L_l^c - f_{Dl}] d\Gamma_D(c) \right\} \right) = f_D^e,$$

where τ_{Dl} and f_{Dl} represent the export costs from D to $l = H, F$; L_l^c and λ_l represent market size and competition in export market l (exogenous due to the ‘small’ country assumption); and \hat{c}_{Dl} represent the export cutoffs. Although those cutoffs are endogenous, they only depend on the exogenous demand conditions in l (along with the exogenous export costs). Thus, those cutoffs do not depend on demand conditions in D (either L_D^c or λ_D). Thus, the average export profits in the condition above can be written in terms of variables that are exogenous to those demand conditions; moving those average profits to the right-hand side, we obtain:

$$\sum_{m=0}^{\infty} \left[\int_0^{\hat{c}_{DD}/z_D(m)} [\pi(cz_D(m), \lambda_D) L_D^c - f_{DD}] d\Gamma_D(c) \right] = f_D^e - \sum_{l=F,H} \left(\sum_{m=0}^{\infty} \left\{ \int_0^{\hat{c}_{Dl}/z_D(m)} [\pi(\tau_{Dl}cz(m), \lambda_l) L_l^c - f_{Dl}] d\Gamma_D(c) \right\} \right).$$

The left-hand side of this free entry condition is identical to the one we derived in (15) for the case where there are no exports from D . Since the right-hand side does not depend on market conditions

in D , all of our comparative statics with respect to those conditions will remain unchanged (it is as-if we changed the level of the entry cost for D to this new right-hand side level).

C Robustness Tables

Table 14: Impact of Purged Trade Shock on Firm Productivity

Dependent Variable Specification	log prod.		Δ log prod.		log prod.		Δ log prod.	
	FE	FD	FD-FE	FE	FD	FD-FE		
log (trade shock \times export intens.)–purged	0.118 ^a (0.021)			0.101 ^a (0.021)				
$\tilde{\Delta}$ (trade shock \times export intens.)–purged		0.128 ^a (0.023)	0.134 ^a (0.027)				0.136 ^a (0.024)	0.141 ^a (0.028)
log capital stock per worker				0.117 ^a (0.004)				
log raw materials				0.086 ^a (0.003)				
Δ log capital stock per worker							0.125 ^a (0.005)	0.133 ^a (0.006)
Δ log raw materials							0.092 ^a (0.003)	0.090 ^a (0.003)
Observations	213001	185688	185688	203977	175619	175619		

FE refers to firm-level fixed effects. All regressions also include industry-year dummies. Standard errors (clustered at the firm level) in parentheses: ^c < 0.1, ^b < 0.05, ^a < 0.01.