The Effects of Exporting on Labour Productivity: Evidence from German Firms

Johannes Schwarzer

March 2017

CEP Working Paper 2017/2

About the Author

Johannes Schwarzer is a fellow at CEP with a focus on the links between services trade and sustainability. He is also a PhD Candidate in Economics and Finance at the University of St.Gallen and holds a Masters Degree as a MacJannet Scholar from the Fletcher School / Tufts University in Boston and a Masters Degree in International Relations from the Graduate Institute of International and Development Studies in Geneva. Prior to joining CEP, Mr. Schwarzer has been part of the management team of oikos as a PhD fellow and has worked on assignments with the International Institute for Sustainable Development (IISD) and the World Trade Organization (WTO).

The CEP Working Paper Series is published by the Council on Economic Policies (CEP) — an international economic policy think tank for sustainability focused on fiscal, monetary, and trade policy. Publications in this series are subject to a prepublication peer review to ensure analytical quality. The views expressed in the papers are those of the authors and do not necessarily represent those of CEP or its board. Working papers describe research in progress by the authors and are published to elicit comments and further debate. For additional information about this series or to submit a paper please contact Council on Economic Policies, Seefeldstrasse 60, 8008 Zurich, Switzerland, Phone: +41 44 252 3300, www.cepweb.org – info@cepweb.org

Abstract

We revisit the "self-selection vs. learning-by-exporting (LBE)" debate with new evidence on a large panel of German firms of all economic sectors up to the 3-digit NACE level, between 1993-2014, and shed new light on the channels that foster export-induced productivity gains. In line with previous results, we find substantial pre-export differences in productivity between future exporters and domestic firms. Nevertheless, these pre-export differences remain constant over time and we find strong evidence *against* a conscious self-selection effect, in which firms would actively engage in increasing their productivity in temporal proximity to starting to export. In contrast, we find strong support for the learning-by-exporting hypothesis in both the manufacturing and the services sector. However, the learning effect is not progressive and more short-lived in the latter than in the former. We explain the different sectoral performances with significant differences in access to foreign markets, which is substantially lower and more concentrated within few firms in services. Furthermore, we show that across sectors, the size of the LBE effect depends on the level of within-sector competition. In line with basic microeconomic theory, productivity gains are higher for entrants into exporting, which operate in relatively uncompetitive domestic sectors, pointing to an important competitiveness channel for productivity gains. Our results suggest that the services sector offers the largest scope for productivity gains through trade policies aiming at facilitating market access.

Acknowledgements

I would like to thank Andrea Ariu, Alexander Barkawi, Simon Evenett, Christian Viegelahn, and seminar participants at the second EUI *Empirical Investigation in Services Trade* Conference in Florence for excellent comments on an earlier draft. This study uses the IAB Establishment Panel, Waves 1993 - 2014. Data access was provided remotely within project fdz1103 via JoSuA at the Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB). I would like to thank the collaborators of the Centre for their excellent assistance. All remaining errors are my own.

Contents

1 Introduction

International trade has come under heavy fire notably in the public debate in developed economies since the global shift in manufacturing (most prominently in Autor et al. (2013)). Two focal points of that debate are remarkable: 1) The focus on increased import competition, and 2) the focus on the manufacturing sector. In this paper, we switch the focus to the exporting side and broaden it over the entire range of economic activity, revisiting and deepening an issue that merits further investigation: the productivity effects of exporting.

In the workhorse theory of international trade as developed by Melitz (2003), firms are endowed with different productivity draws, which are predetermined and unchanging over time. Only those firms obtaining a productivity draw above the threshold for exporting will enter foreign markets. Indeed, there is widespread empirical backing for this prediction that firms engaging in exports are on average more productive than their purely domestically operating counterparts (see e.g. Bernard et al. (2012) for a recent literature review).

In reality, of course, firm productivity levels may be endogenous to firm decisions and may hence also change over time. Similarly, entry and exit into exporting are recurring features of individual firms. Over recent years, there has been increased interest in better grasping the direction of causality in the strong correlation between productivity and exporting.

Two hypotheses are generally put forward to explain the mechanism underlying the "black box" of higher observed productivity in exporting firms: Self-Selection and Learning-By-Exporting (LBE).

Self-selection into exporting implies that firms with higher productivity "self-select" into exporting, as their productivity edge allows them to amortize the higher costs of serving foreign markets. Forward looking firms will in this view increase productivity in preparation for tapping into foreign markets. A broad range of empirical studies reviewed in Wagner (2007), Greenaway and Kneller (2007), and Bernard et al. (2012) confirms substantial differences in firm-level productivity between domestically operating firms and future exporters prior to their entry into exporting. The *Learning-By-Exporting (LBE)* hypothesis stipulates that firms increase their productivity as a consequence of exporting. According to this hypothesis, the productivity-increasing effect of a firm's international activity is a consequence of e.g. increased competition stemming from the larger international market and knowledge and expertise related to the foreign market, which non-exporters do not possess. The evidence for this effect so far is rather sparse (see e.g. Hosono et al. (2015), Manjón et al. (2013), Lileeva and Trefler (2010), De Loecker (2007), Van Biesebroeck (2005)).

Part of the paucity of evidence for the latter may be explained by misspecifications of the estimation process for productivity. Most studies estimate total factor productivity through a variety of methodologies, as reviewed in Van Beveren (2012). However, none of these estimation techniques take a potential productivity effect of exporting into account. A subsequent (matched) difference-in-difference analysis between (future) exporters and domestic firms - as is standard in this literature - will hence eventually translate into a bias against the LBE hypothesis, as pointed out by De Loecker (2013). This shortcoming in the existing literature alone warrants further country-level analyses.

While focusing only on labor productivity, our study uses a robust set of specifications to investigate the LBE effect in Germany. Our study will investigate the effects of exporting on labor productivity for firms of all economic sectors. As noted by (Wagner (2012), p.23), while "we have evidence on the links between international trade and productivity in manufacturing firms from a large number of empirical studies published during the past 15 years, comparable information for firms from services industries is scarce and of a recent vintage". General comparability of firm characteristics in the context of international trade in goods and services was first confirmed by Breinlich and Criscuolo (2011) on a large sample of UK firms. Vogel and Wagner (2011) find a statistically significant exporter premium for firms in German business services sectors (NACE 72, 73, and 74) between 2003 and 2007. However, this premium appears to be driven by outliers and becomes insignificant once they control for those in their regression. For the same time period, sectors and comparing German data with available data from France and the UK, Temouri et al. (2013) find no evidence for LBE for various measures of firm performance. Using a very comprehensive dataset on Danish firms in services and manufacturing, Malchow-Møller et al. (2015) are able to disentangle services and goods traders and investigate the respective links with long term (2002 - 2008) productivity growth. Their findings suggest that firms that have started exporting goods in this period have experienced higher average productivity growth than firms that have never exported in this period. Having started to export services is also associated with increases in productivity growth, but less so and only for firms in the services sector.

In this paper, we contribute to the literature on the productivity effects of exporting by proposing an unprecedented look at productivity developments in temporal proximity of *each firm's* first entry into exporting. We are exploiting variation in productivity per firm, which allows us to control for inherent differences between types of firms that standard difference-in-differences estimation is not able to detect and matching approaches can only approximate for. Unlike previous studies that have focused on firm level determinants of productivity gains from LBE, we provide evidence for sector-level determinants that help explaining both the existence and the magnitude of the LBE effect. To this end, we are using a large panel of German firms spanning the period from 1993-2014, exploiting the panel structure to identify a causal effect and disaggregating our analysis up to the 3-digit NACE level, hence comprising both exports in the manufacturing and services sector.

In line with previous results, we find substantial pre-export differences in productivity between future exporters and domestic firms, across all sectors, but indications for less important differences in the services sector. Nevertheless, these differences remain constant over time and we find strong evidence *against* a conscious self-selection effect, in which firms would actively engage in increasing their productivity in temporal proximity to starting to export. In contrast, we find strong support for the LBE hypothesis in both the manufacturing and the services sector, as average productivity rises after initial entry into exporting, regardless of whether the export status is maintained in subsequent years or not. However, the effect is stronger in manufacturing firms than in services firms. The former exhibit increasing yearly productivity growth rates even more than two years after exporting, while the productivity growth rates of the latter group decrease (albeit remaining above pre-export averages). We explain the different performances of the manufacturing and services sector with significant differences in foreign market access and propensities to export and are able to show that across sectors, the size of the LBE effect depends on the level of domestic within-sector competition.

The next section will describe the dataset this study uses and discusses the choice and calculation of variables. Section 3 contains the analysis, which we split into a preliminary analysis to check for broad comparability with other studies in 3.1, and an extended analysis that contains our main methodology and results. We will conduct robustness checks in section 4 and conclude in section 5.

2 DATA

We use confidential, representative German establishment-level data[1], which is managed by and kindly provided through the Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB) (Ellguth et al. (2014) and Fischer et al. (2009)). The database offers a wealth of information to test the above-mentioned hypotheses. A firm level unique identifier allows us to observe firms over time and we hence link each wave of the survey (roughly 14.000 yearly responses) over time to obtain a panel on key firm characteristics for the period 1993-2014. As responses are not always complete, firms enter and exit the survey, and the set of variables that is asked is not always constant, we obtain a very large unbalanced panel with key observations on total turnover, the share of foreign sales in total sales, input volumes, average wages, employment and investment. We list the summary statistics of the main variables used in table (1). Since not all firms answer all questions, we end up with diverging numbers of total observations for each variable. Relatively low feedback on external inputs used result in the relatively low number of observations for our productivity variable which, however, still yields a largely sufficient number (156189) for our purposes. We will describe the construction of these variables in detail in the next section. The dataset has been extensively used for German labor market research, but surprisingly little in trade. As any researcher who is entitled to use the data is contractually obliged to register his publications in a database managed by the Institute for Employment Research, we can easily verify that this paper is the first of its kind.[2]

In comparison with other datasets on German firms, the IAB panel is the most comprehensive, and hence most suitable for our purposes. The research data center of the German Federal Statistical Office maintains a similar dataset (the "AFiD-Panel Strukturerhebung im Dienstleistungsbereich" (FDZ (2015))), but it is restricted in scope (only certain services industries), time (2003-2010, with a methodological break in 2007/2008) and includes only firms with a turnover of more than 250.000 Euros per year. The German Federal Statistical Office also maintains tax records for the universe of German firms, which can be accessed for research purposes. While this dataset is perhaps the most accurate and comprehensive of all, it only

[1] For ease of exposition, we will henceforth refer to establishments as firms.

[2] In fact, Vogel (2011) does examine a very similar question, but bases his analysis primarily on data from the German Federal Statistical Office, and the data from the IAB is used only for robustness checks. His analysis differs in methodology, is narrower in scope (business services) and uses only years 2000-2005.

Table 1: Data description

Variable	Obs	Mean	Std. Dev.	Min.	Max.
exp	213931	.2298171	.4207161	0	1
starter	213931	.2923559	.454846	0	1
log productivity	156189	10.48702	.9910139	-.0497428	16.84333
log employment	213931	3.163625	1.697547	.6931472	10.97972
log investment	200839	4.438881	4.133074	0	16.65125
log wages	188037	7.086665	1.290652	0	10.59666
log dom sales	180886	11.17106	1.112905	-.0612103	18.35064

records exports of goods, as services exports are not tax-exempt. In contrast, the German Bundesbank maintains a very detailed record of all international services transactions of all German firms, tracking detailed industry affiliation, as well as type of service transaction and destination country over time. However, no other firm level characteristics are provided and the law prohibits matching this data with either datasets from the IAB or the German Federal Statistics Office.

Our dataset only contains information on the share of foreign sales in total sales. As such, we cannot distinguish between services and goods trade, such as e.g. in Malchow-Møller et al. (2015). This can be both an advantage and a disadvantage. On the one side, for analytical purposes it would be illuminating to have a better grasp on the type of export that a certain firm in a certain sector is associated with. For instance, we know that services trade, goods trade, the manufacturing sector, and the services sector are closely intertwined. For Germany, we know that manufacturing firms account for almost 25% of services exports (Kelle and Kleinert (2010)). Conversely, 14% of services firms in Denmark appear to be exporting goods (Malchow-Møller et al. (2015)). On the other hand, it is increasingly complex to disentangle goods and services in general, as manufacturing firms both buy and produce more services in-house than before, but also sell and export more services than before (Lodefalk (2013)). Indeed, it appears that the services content of international trade in goods appears to have been systematically underestimated until recently (Cernat and Kutlina-Dimitrova (2014)). Our data hence take a rather agnostic approach toward the exact type of international transaction, but the fact that our panel is on an establishment level may therefore actually add confidence to associating sectoral exports with the corresponding type of export.

Table 2: Sectors and NACE Codes

Sector	Rev 1.1	Sector	Rev 1.1
Agriculture, Hunting, Fisheries	1,2,5	Telecommunication	643
Mining & Quarrying	10,11,12,13,14	Transport, travel & storage	60, 61, 62, 63, 641
Food Products, Beverages & Tobacco	15, 16	Finance & Insurance	65, 66, 67
Textile & Leather	17, 18,19	Real Estate	70
Wood, Paper & printing	20,21	Renting	71
Coke & Refined petroleum products	23	R&D	73
Chemical, Pharmaceutics	24	Legal, Accouting, Consulting & advertising	744, 741
Rubber, Plastic & Non-Metallic Minerals	25,26	Architecture & Engineering	742, 743
Basic & Fabricated metals	27, 28	Other professional, scientific or technical services	748
Machinery	29	Employment, Security & Investigation,	745, 746, 747
Computer, Electronic & Optical	30, 32, 33	Public Admin, Defense, Social Security	751, 752, 753
Electrical Equipment	31	Education	80
Motor Vehicles & other Transport equipment	34, 35	Health	851
Furniture, Sport Goods, Toys, & other	36	Veterinary	852
Utilities	37, 40, 41, 90	Social Services	853
Construction	45	Art, Entertainment & Recreation	923, 925, 926, 927
Trade & Repair	50, 51, 52	Other Services	93
Hotels and Restaurants	55	Households	95
Audiovisual Media and Broadcasting	22, 921, 922	Extra-Territorial Organizations	99
IT Services	72, 924	Unclassifed	N/A

2.1 Industry Classifications

During the period of observation, the system of industrial classifications has undergone two changes, NACE Rev. 1.1 in 2003 and NACE Rev. 2 in 2008. In order to obtain time-consistent classifications of industry codes, we merge our dataset with correlation tables obtained from Eberle et al. (2011). Their identification strategy for the generation of time-consistent industry codes basically comes from the fact that in the years of conversion firms were required to indicate both their new and their old industry codes. We chose NACE 1.1 as our reference code and hence obtain time-consistent 5-digit codes, which we aggregate into the classification displayed in table (2).

We also do not observe an industry classification for firms before the year 2000, except for a self-reported more general branch affiliation (industry classifications are otherwise assigned based on administrative records). Here, we make the assumption that firms that are also observed in earlier years belong to the same industry classification they belong to in 2000 and fill in the unobserved data accordingly.

2.2 Rescaling The Time Variable

We have a large unbalanced panel that ranges from 1993 to 2014, with firms dropping in and out of the sample for reasons we cannot observe. We do not know whether the firm has ceased to exist, or whether it simply discontinued answering the survey. By design, the survey is organized in such a way as to put significant resources into getting the same establishments to respond that were responding in the previous year, as witnessed by the large number of observations listed in table (1), listing - by construction - only observations on firms that appear at least in two consecutive years: As we want to observe firms over time, we delete all single observations from the sample. Furthermore, in order to be able to assess the effect of exporting, we need to find a common scale to all firms. We hence create a time variable that counts the intervals in years from the moment a firm is first observed to export, which we denote as zero. A firm that is observed to export from the beginning hence appears only for time intervals > 0, counting the remaining years of observing that particular firm in the sample. On the other hand, a firm that is not initially an exporter will be observed for the time intervals < 0 until it is observed to export for the first time (at $t = 0$), and *all remaining* $t > 0$ years. For strictly domestically operating firms that never start to export, the value zero is simply assigned to their rounded mean period of observation, which is expected to be random across firms in the sense that there is no particular or systematic importance to that moment of time.

2.3 Creating the $STARTER$ variable

Armed with the rescaled time variable, we need to decide on an appropriate measurement for denoting a firm that has started to export. While this undertaking seems to be trivial, it is worth pondering its importance for a while. The early studies on the topic have usually looked at a dummy variable EXP_{it}, which indicate whether industry or firm i has been observed to be an exporter in year t. A voluminous literature has since found support for the resulting finding of large exporter premia, in terms of productivity, average wages, size of the workforce etc. For our purposes, such a dummy indicating the moment a firm exports is not sufficient, as it would give us information only for the years that the firm is observed to actually export. Another frequently used indicator is a dummy variable that takes on the value of 1 if a firm exports in a given year, but has not been exporting in the previous year. Again, we believe that such a variable is not sufficient for our purposes, as a firm may well be classified as starting to export several times during the time it is observed.

We want to test whether a firm "learns" from exporting, i.e. whether we can observe any significant change of the dependent variable in response to a single change in the independent variable. In order to hence characterize each firm by a single metric, we generate the dummy variable $STARTER_{it}$ that takes on the value of 0 if a firm is never observed to export throughout its appearance in the panel, 1 for firms that export throughout all observations in the dataset. For firms that are initially observed not to export, their value of $STARTER_{it}$ is zero until the moment they first export (at time $t = 0$), and 1 for the remaining observations, regardless of its export status. We hence have three groups of firms in our sample: (1) domestic firms that never export, (2) switchers, for whom the value of $STARTER_{it}$ switches to 1 after the first time they have been observed to export, and regardless of whether they keep exporting or not, and (3) international firms that have been observed to export throughout their observations.

Of course we do not observe whether a firm has been exporting before it has been included in the survey. This shortcoming may potentially create a pro-LBE bias, as a firm that enters the survey at a time where it does (coincidentally) not export may display higher productivity than a firm that has never been exporting before. However, we neither observe whether a firm will start exporting after having opted out of the survey (for reasons we do not observe), creating a potential bias in the opposite direction. Since the design of the IAB survey does not pay attention to export status when selecting the firms in the sample, we assume that these two effects should cancel each other out, or at least reduce the significance of such a potential bias. In addition to mere faith in such randomness, however, we exploit differences across groups and time in within-firm variation to isolate the net effect of exporting, correcting for any potential time invariant biases on the firm level. We hence believe that this classification best fits our purposes.

2.4 How to measure productivity?

In the literature, productivity is often estimated building on the methodology developed by Pakes and Olley (1995), which consists of estimating a production function in a first step in order to get to a measure of total factor productivity (TFP). In fact, several ways exist to estimate TFP, each having specific properties that may be useful in different contexts, but appear to yield very similar results across methodologies (Van Beveren (2012)). Unfortunately, values for the capital stock are not reported in our data. Since we do observe investment levels, we might obtain a measure of the capital stock through applying the perpetual inventory method. However, given sometimes patchy investment data and short time spans of firm observations, we are doubtful of whether this approach would add value (see also the discussion in Müller (2010)). Instead, following Lileeva and Trefler (2010), we use a different measure of labor productivity, constructing our variable as value-added per worker. In constructing our productivity variable, we use the total number of workers. This choice is motivated by the fact that the shares of other available variables such as high-skilled, temporary and short-term employment remain remarkably invariant over intervals of observation t within the three groups.

3 Analysis

3.1 Pooled OLS

We now turn to a preliminary analysis of our dataset. In order to ensure comparability with similar studies, we follow Manjón et al. (2013), De Loecker (2007), and Bernard and Jensen (1999), estimating variants of the following equation by Ordinary Least Squares:

$$prod_{ikt} = \alpha + \beta I_{ikt} + \gamma l_{ikt} + \delta_{year} + \lambda_k + \varepsilon_{ikt} \quad (1)$$

$prod_{ikt}$ refers to the log of labor productivity of firm i in industry k at time t. The I_{ikt} variable is replaced by EXP_{ikt} in a first step and then by our $STARTER_{ikt}$ variable and l_{ikt} is the log of employment. Our coefficient of interest is β. Given the rescaling of our time variable, we use calendar year dummies to keep track of year specific effects δ_t (such as business cycle or other year-specific shocks), and finally λ_k is an industry specific fixed effects that controls for differential, time-invariant productivity tendencies across industries as classified in (2). Given the log specifiation of equation (1), we can interpret $(exp(\beta)-1)*100$ as the percentage difference between firms for which $I_{ikt} = 1$ and those for which $I_{ikt} = 0$.

Table 3: Pooled OLS

	1a	1b	2a	2b	3a	3b	4a	4b
exp	0.493	0.327						
	(0.000)	(0.000)						
starter			0.479	0.333	0.390	0.297	0.170	0.117
			(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
labor	NO	YES	NO	YES	NO	YES	NO	YES
N	156189	156189	156189	156189	132039	132039	32716	32716
R^2	0.152	0.181	0.156	0.184	0.131	0.154	0.069	0.104

Columns (a) refer to regressions without the employment control variable, (b) columns refer to those with the employment variable, p-vales in parentheses

We first estimate by pooled OLS, with both fixed effects and without the employment variable and display the results in the (a) columns of table (3). The first column tells us that exporting at any given year is associated with an average of roughly 63% higher productivity over the entire sample of firms. This figure is higher than the usual roughly 35% productivity premia the literature generally finds - mainly, because we do not control for firm level log employment. The moment we do so, displayed in columns (b) of table (3) - we find a highly comparable number of roughly 38% productivity premium for exporters.

In a next step, we regress equation (1) again over the entire sample of firms, but using the $STARTER$ variable as the dependent variable, where - interestingly - the coefficient does not change much with respect to the first specification.

We still have no idea whether the productivity differences are inherent to the firm, or whether they are associated with the specific exporting status we have defined in the $STARTER$ variable. We suspect that international firms may bias our coefficient upwards as they may be more productive in the first place. In order to get a better grasp on this question, we perform the same set of regressions on a restricted sample, which only includes domestic and switching firms (columns (3) of table (3)). As suspected, the coefficient decreases slightly in magnitude: Switching firms after their first observed year of exporting are still roughly 35% more productive than the average of domestic firms and switching firms before exporting, after controlling for firm-level employment. This result implies that we can now rule out large productivity differences between international firms and firms that have started to export. Nevertheless, we still do not know whether there are significant pre-export differences between domestic firms and switchers. Column (4) displays the results of the same set of regressions on the sub-sample of switching firms only; we hence compare the mean productivity of switching firms only before and after exporting. Controlling for employment, we find that the difference is still significant, but less than half as important as in the previous set of regression, suggesting substantial mean pre-export productivity differences with domestic firms.

3.2 Comparing Means

In the previous section, we have established significant mean differences between domestic firms and both switchers, as well as international firms. We have also found significant differences among switchers before and after exporting, amounting to an average percentage difference of roughly 12%. We are still unsure as to how to interpret these results in the light of the self-selection, as well as the LBE hypothesis. Substantial pre-export differences between domestic firms and switchers suggest that self-selection is certainly at play, but how do we interpret the fact that post-export average productivity is even higher among switchers? Before resorting to more sophisticated econometric techniques to shed more light on these questions, we proceed with a simple graphical analysis. We compute

Figure 1: Comparing Means: Labor Productivity

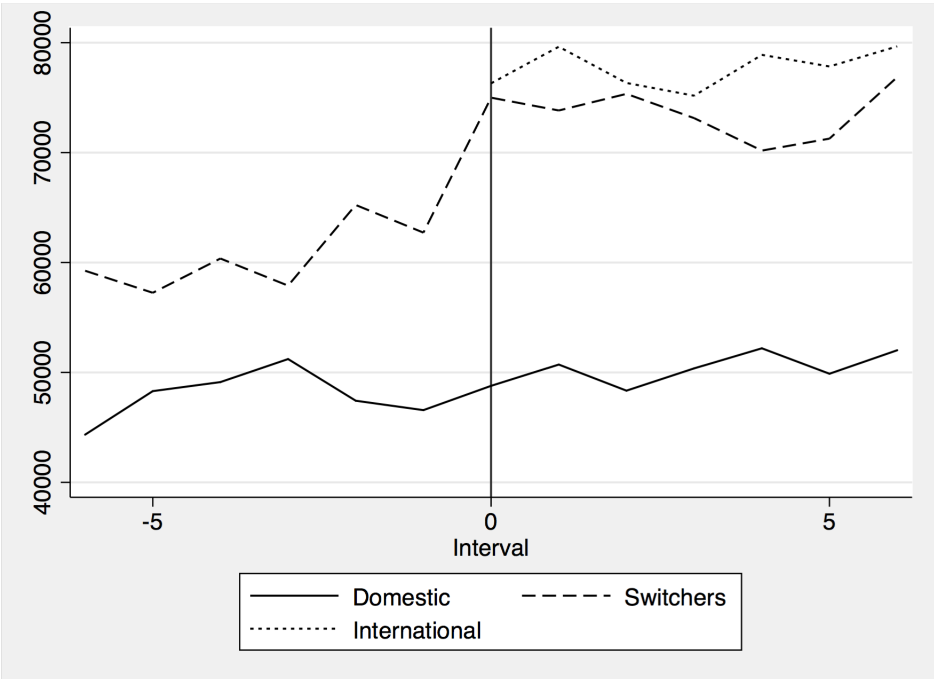

the mean productivity levels of each of our three groups of firms for each interval of observation and plot the results in figure (1). The graph very nicely reflects our regression results, but also gives illuminating insights on the phenomenon we want to explain. We indeed observe substantial pre-export productivity differences between domestic firms and switchers, but these appear to be relatively constant. The graph seems to suggest that relatively more productive firms do self-select into exporting, but not in that they increase their productivity in temporal proximity to their entry into export markets. If we drew a trend line for pre-export observations of both groups, they would both be quite flat. It is only once exporting has occurred (interval $t = 0$) that average productivity increases, to levels comparable with international firms.

Figure 2: Comparing Means: Employment

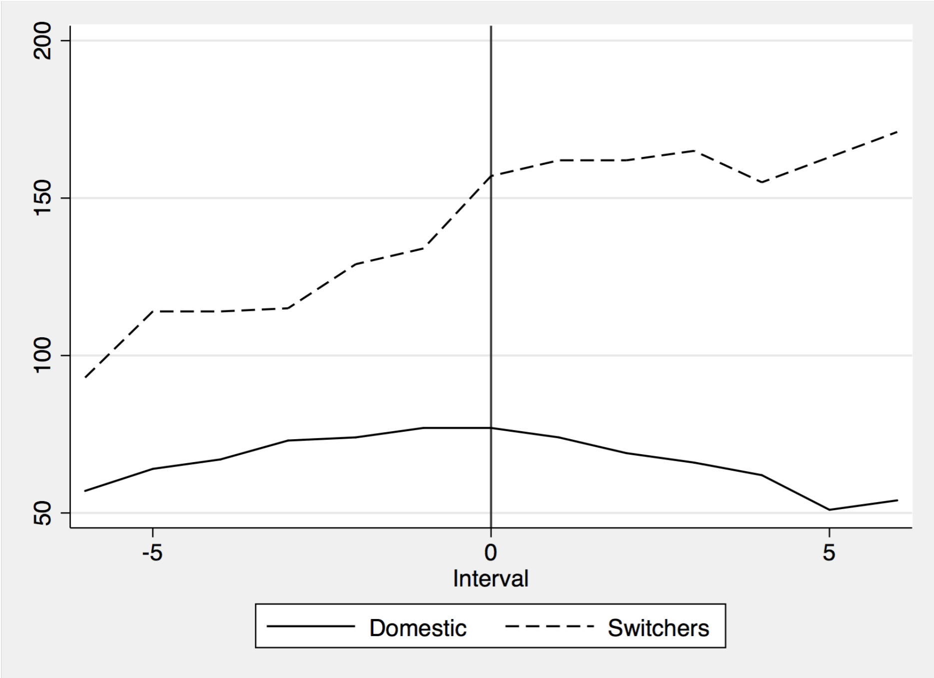

Figure 3: Comparing Means: Inputs per worker

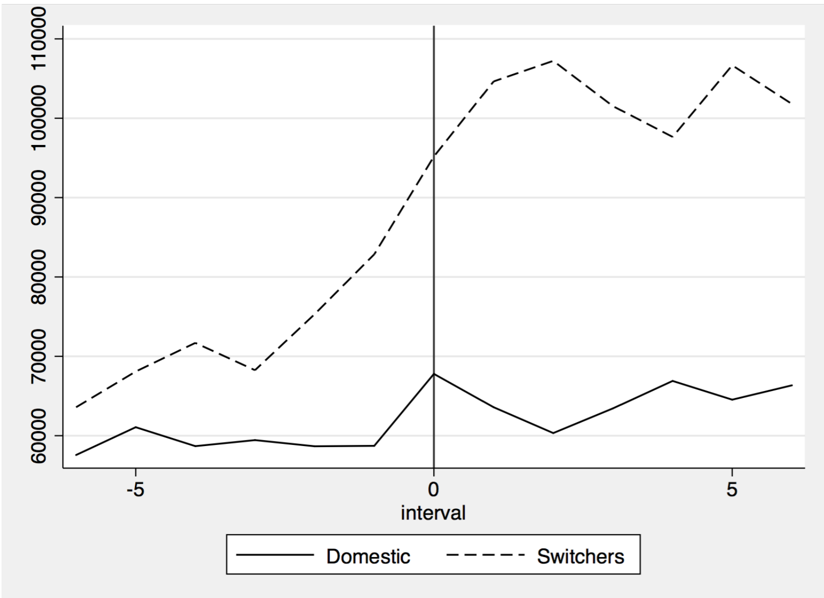

Opening a small parenthesis, it is instructive to examine labor productivity jointly with labor and external inputs purchased by the firm. We compute the means for employment and inputs and plot them in figures (2) and (3).[3] It is interesting to note that the trends for average employment of switching firms is increasing throughout and at similar rates, while the average size of domestic firms does not display any particular trend, except for being smaller towards the tails. Seen in the light of rising employment throughout, the labor productivity increases that occur with exporting appear to be even more spectacular and not to be driven by reducing the average workforce. The stark average increase in purchased external inputs confirms larger demand for those per worker. Unreported figures for investment per worker draw a similar picture. Both metrics, however, do not display a substantial pre-exporting jump, which implies that firms seem not to make conscious pre-export choices concerning the volume of these metrics.

3.3 Pooled OLS with firm fixed effects

Our analysis so far has shown that there appear to be intrinsic differences between our groups of firms, and we cannot say much about those unless we control for more firm specific effects. We begin to do so by replicating the regressions in table (3a), without any firm level controls but with a year fixed effect and this time a firm fixed effect, as in equation (2).[4]

$$prod_{it} = \alpha + \beta I_{it} + \delta_{year} + \lambda_i + \varepsilon_{it} \tag{2}$$

If there are indeed intrinsic, firm-specific and time-invariant differences, the λ will pick those up and β will provide us with a more accurate estimation of the percentage difference each regression aims at uncovering. A quick look at the results in table (4) confirms this intuition, notably the high R^2 we obtain without any firm-level covariates except the dummy variable. We follow the same procedure as before, where column (1) and (2) display the results of a regression over the entire sample, including international firms. The coefficients are again very similar and highly statistically significant, but much lower in value. Controlling for firm fixed effects (such as intrinsic productivity differences), we find that exporters and starters are now just 2.7% more productive than non-exporters and non-starters as we defined them. Interestingly, in column (3) we observe that removing international firms

[3] We do not include international firms here, as they are of significantly larger average size and would make the graph less readable.

[4] An unreported Hausmann test confirms the appropriateness of fixed effects over random effects and the necessity for year fixed effects

Table 4: Pooled OLS with firm fixed effects

	(1)	(2)	(3)	(4)
exp	0.027			
	(0.0011)			
starter		0.027	0.035	0.036
		(0.0096)	(0.0013)	(0.0035)
N	156189	156189	132039	32716
R^2	0.727	0.727	0.715	0.651

p-values in parentheses

yields a higher coefficient β, reinforcing the LBE hypothesis in that it underlines the relevance of starting to export for productivity gains. The same is true for the value of β obtained in a regression over switching firms (column 4), which tells us that switching firms are on average 3.7% more productive once they have started to export, controlling for their individual average pre-export productivity levels.

3.4 Fixed-effect estimation and sectoral decomoposition

We refine our analysis further by focusing solely on within-firm variation, using a fixed effect estimator. This necessitates the addition of firm-level covariates that are reasonable for our purposes. We hence proceed to estimate a model of the following form:

$$prod_{it} = \alpha + \beta STARTER_{it} + \gamma \mathbf{X}_{it} + \pi_{year} + \lambda_i + \varepsilon_{it} \tag{3}$$

, where $\mathbf{X}_{it}$ indicates a set of firm-level covariates. We opt to control for firm size by including the log of employment, the log of investment per worker, as well as the payroll per worker. Additionally, we include the log of domestic sales per worker, in order to better isolate the effect of exporting, controlling for the purely domestic sources of productivity gains that may occur to firms regardless of their exporting status.

In a panel setting like ours, idiosyncratic errors are likely to be serially correlated. Bertrand et al. (2004) show that the usual standard errors of the fixed effects estimator are drastically under-estimated in the presence of serial correlation. As suggested by Stock and Watson (2008), we cluster standard errors on the firm level to control for both heteroskedasticity as well as within-firm serial correlation.

Finally, we group firms into a manufacturing and a services sector and proceed within these groups as above, regressing over a) all three types of firms (domestic, switchers and international), b) only domestic and switchers, c) only switchers. We plot the results in table (5).

Table 5: Fixed Effects Estimation

	All			Manufacturing			Services		
	a	b	c	a	b	c	a	b	c
starter	0.087	0.107	0.088	0.112	0.145	0.111	0.058	0.070	0.078
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0110)	(0.0027)	(0.0029)
firm controls	yes								
firm fe	yes								
year fe	yes								
N	142448	121025	29896	39170	24823	11287	78135	73676	14330
R^2	0.157	0.208	0.102	0.103	0.145	0.081	0.164	0.200	0.110

Manufacturing comprises sectors 3-14, Services comprises 16-37. Standard errors are clustered on the firm level and p-values are given in parentheses

Qualitatively, the results are similar to what we have established so far, except that we are now looking at within-variation only, which enables us to get rid of self-selection effects that may occur as the result of inherent differences in firm productivities. Quantitatively, the coefficients we estimate are larger. Looking at column (b) of the regression over all firms, we find that switchers are on average almost 11% more productive once they export, compared to domestic firms and before exporting. Looking at switchers only before and after starting to export, we find a significant starting premium of over 9%, which is both statistically and economically highly significant.

Furthermore, we can now for the first time look at differences between firms in the manufacturing and in the services sector. Overall, sectoral results resemble the aggregate results. Starting to export is associated with higher productivity gains in manufacturing than in services, but the effect is statistically and economically significant in both sectors. Compared to the aggregate analysis and the one on manufacturing, the analysis on services firms displays an interesting peculiarity: The coefficient in column (c) is higher than in column (b), suggesting that the productivity differences between domestic firms and switchers are less substantial than in the manufacturing sector.

Still, these are large economic sectors and we hence dig deeper into detailed industry classifications to get an idea of which sectors are those where starting to export is associated most closely with productivity gains. To this end, we estimate equation (3) over firms in each subset of industry classifications as generated in table (2). We find that indeed not all industries seem to be associated with LBE effects. We list those industries where we find a statistically significant coefficient β in table (6).

Table 6: LBE Industries

LBE Manufacturing	LBE Services
Wood, paper and printing	Construction
Chemical and pharmaceutical products	IT services
Rubber, Plastic and non-metallic mineral products	R&D
Basic metals and fabricated metal products	Architecture and engineering
Electrical Equipment	Other professional, scientific or technical services
Furniture, jewellery, sport goods, toys, and other	Education
	Transport, Travel and Storage
	Health
	Art, Entertainment and Recreation
	Real Estate

Intuitively, we fail to detect an immediate reason for why these precise sectors display the LBE effects we find. While this question is beyond the immediate scope of this paper, we nevertheless ponder it for a moment, to the extent that the limitations we face in our dataset allow us to do so. In fact, firm-level workforce characteristics we have not yet accounted for seem not to be important determinants of these different behaviours, even those that vary across time and, hence, are not picked up by the individual fixed effects employed in our regressions. Importantly, hiring and firing decisions, the share of qualified workers, part-time and temporary employment in total employment are not significantly associated with post-export productivity increases. We have seen in 3.1.2. that management decisions such as investment or purchase of external inputs per worker are significantly associated with productivity increases, comparing switching firms with domestic ones across sectors. However, LBE sectors do not display significant average differences with non-LBE sectors along those lines.

Table 7: Comparing sectors

	Manufacturing		Services	
	Non-LBE	LBE	Non-LBE	LBE
Export Propensity	50.3%	53.4%	11.6%	12.8%
Market Concentration	0.059	0.074	0.068	0.072
Export Concentration	0.086	0.116	0.233	0.281

Theoretically, we expect to find such effects primarily in relatively export-oriented sectors. The average propensity to export is very heterogeneous across sectors, which is in part a reflection of differences in intrinsic exportability of certain goods or services over others. For example, the services sector has long been regarded as non-tradable as a whole. It is through revolutions in technology and transport that this sector is getting increasing attention in the international trade literature. Calculating the potential tradability of different services sectors in the US on grounds of their geographic concentration, Jensen et al. (2011) obtains a ranking of these sectors according to their 'tradability'. While our sectors listed in table(6) are much more aggregated (in an effort to ensure time-consistent classifications of economic activity, as well as an adequate trade-off between sectoral precision and meaningful numbers of observations within these sectors), there is a striking overlap with the sectors identified by Jensen.

Likewise, there are also differences in propensity to export across manufacturing sectors. These can result from a whole variety of factors, ranging from traditional explanations of comparative advantage to differences in consumer valuation of some goods over others. In both cases, our data confirm the heterogeneity across sectors in export propensity.

Table (7) shows that (i) the average propensity to export is, as expected, much lower in services than in manufacturing. While export propensities of above 11% are quite low, it is certainly not the case that the services sector *per se* is not tradable, but scope for exporting is on average much lower than in manufacturing, which in turn may be part of an explanation for lower LBE effects in services as established in table (5). At the same time, table (7) shows that (ii), if we compare LBE and non-LBE sectors within services and manufacturing each, average propensity to export is higher in the former in both cases (roughly 6% in manufacturing and 10% in services), reinforcing our conjecture that export orientation matters for LBE effects.

Apart from export orientation, and hence the simple capacity to tap into foreign markets, we suspect that the degree of competition matters as well. Increased competition is widely considered as a major driver of firm productivity (see e.g. Aghion et al. (2015)) and we hence expect those productivity gains resulting from exporting to be relatively higher in domestic sectors with relatively low levels of competition. The intuition here follows from basic microeconomic theory, which establishes that firms in uncompetitive markets tend to be relatively unproductive as they face little competition. Entry into exporting hence entails productivity upgrades, as firms operating in a formerly uncompetitive sector find themselves in competition on the world market. In order to investigate this channel further, we thus need a measure of the degree of competition within a sector and choose to compute a normalized Herfindahl index per sector as follows:

$$NH_s = \frac{(H_s - 1)/N_s}{(1 - 1/N_s}$$

, where

$$H_s = \sum_i^{N_s} \left(\frac{Rev_{is}}{Rev_s} \right)^2 .$$

The index ranges from 0 to 1, where a higher index indicates higher market concentration. Plotting the results in (7), we do indeed find evidence for higher market concentration, and hence less competition, in LBE sectors, as compared to non-LBE sectors. Again, the difference is less pronounced in the services sector, but these results are suggestive of taking the analysis a step further.

In figure (4) we plot the Herfindahl index we obtain for each sector against the (statistically significant)[5] coefficients we obtain from the individual regression results obtained in table(6). The low number of observations notwithstanding, we have strong suggestive evidence in favor of the hypothesis that LBE effects increase with uncompetitive market structures as proxied by the Herfindahl index. The underlying OLS regression of each sector's β coefficient on its Herfindahl measure yields a coefficient of 0.931, with a 0.260 standard error. A similar hypothesis was put forward by Manjón et al. (2013), who found lower LBE effects for firms in Spain

[5] p-value < 0.05

Figure 4: LBE and Market Concentration

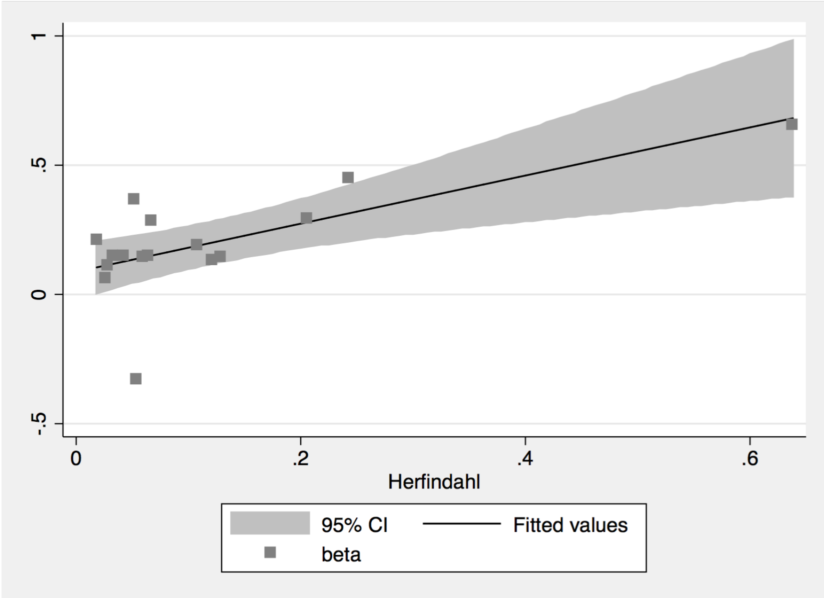

than De Loecker (2007) did for Slovenia, despite using a very similar method. The authors hypothesize that this difference is due to the higher potential of productivity gains in post-Communist Slovenia, without further substantiation, however. Our result is not inconsistent with this hypothesis.

Our result is also robust to the omission of the negative coefficient obtained for "Real Estate". While the sign of this coefficient is somewhat of a puzzle and would require a more in-depth analysis, we believe that the peculiarities of this sector are responsible for the negative association between starting to export and labor productivity. In particular, the real estate sector requires significant local expertise and interaction with clients, which may set it apart from other sectors.

In order to complete our picture, contrasting service sector performance with manufacturing, we also calculate a Herfindahl concentration index for exporting shares only. A higher measure hence indicates the concentration of export revenue in few firms. Unlike the simple measure of export propensity, the concentration index measures the distribution of export revenues among exporting firms. If analyzed jointly with the market concentration index, this metric may point to restrictive access to foreign markets. The last row in table (7) reports the numbers for our sectoral classification. The differences between LBE and non-LBE sectors mimic the differences established earlier between both sectors with respect to market concentration. Intuitively, this result makes sense and has a mechanical component, as export revenue is part of overall market revenue. The higher level of export concentration as compared to market concentration is also readily rationalizeable by the positive correlation between exporting and firm size. What is striking in these results, however, is the sizeable higher concentration in service sector exports, as opposed to its market concentration measures. While the latter are broadly comparable to manufacturing concentration measures, export revenue is highly concentrated in few service sector firms, pointing to highly uncompetitive foreign market access. Not only does the German services sector exhibit a generally lower propensity to export, but even within the group of exporting firms, revenues are highly concentrated.

3.5 Testing for temporal proximity of self-selection vs LBE

Taken together, the previous results suggest that the average post-exporting productivity of switchers is higher than both their average pre-export productivity and domestic firms' average productivity before and after their median observation. This seems to be true for firms in both manufacturing and services industries, where some industries appear to be more predisposed to experience such productivity gains than others. The size of productivity gains on average appears to be higher in the manufacturing sector than in services, which may be the result of differences in the degree of competitiveness of the underlying market structures. We have so far attested self-selection to the extent that future exporters are on average more productive than their domestic counterparts. However, we cannot yet ascertain whether self-selection occurs as an *anticipation effect*, i.e. in temporal proximity to the entry into exporting at time $t = 0$. In other words, we want to test whether exporting granger-causes productivity gains, or whether productivity granger-cause entry into exporting.

In order to do so, we follow the method developed by Autor (2003) and augment equation (3) with leads and lags of the $Starter_{it}$ variable. In particular, we add a dummy for t_{-k}, where k denotes the intervals a firm is observed before entry into exporting, as well as a dummy for t_0 and t_{+j}, where j denotes the intervals a firm is observed after exporting. These dummies each take the value of one only for the year of their corresponding time period and are zero otherwise. We also include a dummy that takes on the value of 1 for all observations $> k$, starting in $t_{+(k+1)}$. Note that all international firms will not enter the sample, since their value of t_{-j} for $j > 0$

is undefined. We therefore regress only over those firms whose time dimensions range from at least -1 to at least +2. We test different values for k and j, with very similar results across specifications. As the number of firms observed drops significantly with larger values of k and j (as we increase the required number of consecutive observations), we display the results of a regression with $k = 1$ and $j = 1$ in table (8), implying that we regress over all firms that are observed at least for a period of four consecutive years. All firms that do not satisfy the criterium of at least one observation prior to exporting (eg international firms), as well as at least two observations after the year of exporting, do not enter the regression. We regress over all firms in column (1), over manufacturing and services firms in columns (2) and (4) respectively, and finally over those subsectors we identified in table (6) as being particularly prone to LBE effects in manufacturing (3) and services (5).

If we would observe an anticipation effect in the sense that a firm makes a conscious effort to upgrade productivity prior to entering into exporting, we would observe a positive coefficient on t_{-1}. The interpretation of that coefficient would be that its productivity at that time exceeds its average productivity when $t_{-1} = 0$, meaning all other years of observation of the firm. In contrast, an LBE effect would be supported by positive coefficients on $t_{\geq 0}$.

Table 8: Leads and Lags

	All	Manufacturing	LBE Manufacturing	Services	LBE Services
t_{-1}	-0.076	-0.058	-0.055	-0.08	-0.078
	(0.001)	(0.073)	(0.249)	(0.000)	(0.062)
t_0	0.108	0.095	0.163	0.116	0.216
	(0.000)	(0.005)	(0.000)	(0.000)	(0.000)
t_{+1}	0.078	0.124	0.166	0.052	0.11
	(0.000)	(0.000)	(0.000)	(0.023)	(0.015)
$t_{+(2)}$	0.056	0.161	0.241	-0.020	0.090
	(0.010)	(0.000)	(0.000)	(0.297)	(0.043)
firm controls			yes		
firm fe			yes		
year fe			yes		
N	101211	18642	9998	63402	29078
R^2	0.253	0.191	0.198	0.242	0.286

Standard errors are clustered on the firm level. p-values in parentheses

Our results in table (8) are remarkably clear, in that the coefficient on t_{-1} is never positive. This gives us confidence that we can reject the null hypothesis that firms self-select into exporting by upgrading their productivity just prior to starting to export. Conversely, we find ample backing for the LBE hypothesis. The coefficients on $t_{\geq 0}$ are very interesting when we compare sectors. The manufacturing sector, and notably LBE manufacturing, displays the predictions of LBE to the letter. At $t = 0$, the average manufacturing firm is almost 10% more productive than its average (18% in LBE manufacturing). At $t = 1$, that firm will be already 13% more productive (18% for LBE manufacturing). For $t \geq 2$, average productivity rises further to 17% (27% in LBE manufacturing). These results suggest that firms literally "learn" from exporting, in terms of productivity gains, as time passes.

In the services sector, the results are not as clear-cut. The coefficients on $t_{\geq 0}$ are also positive, but decrease in magnitude as t rises. For the services sector as a whole, the coefficient on the forward variable $t_{+(k+1)}$ becomes insignificant, whereas it remains significant in the LBE services sector. The same pattern holds when $k = 2$.[6] These results still support the LBE hypothesis, as firms remain more productive than prior to exporting. However, it seems that the learning effect is not progressive and more short-lived than in the manufacturing sector, reflecting underlying differences in competitiveness of market structures as established in 3.2.2.

4 Robustness Checks

Our analysis has consistently focused on average productivity effects. Here we estimate a simple production function to check for differential marginal productivity effects of entry into exporting. We therefore need an employment variable that captures employment before having exported and after having done so for the first time, analogously to our previous analysis. To obtain this, we generate a $nonstarter$ variable that takes the opposite values of our starter variable and is hence 1 for any firm that does not export or has not done so yet, and 0 else. We interact both the $starter$ and the $nonstarter$ variable with firm employment, take logs and estimate the following production function again by fixed-effect estimation:

[6] Results are not reported

$$lnVA_{it} = \alpha + \beta_1 nonstarterl_{it} + \beta_2 starterl_{it} + \gamma cap_{it} + \pi_{year} + \lambda_i + \varepsilon_{it} \quad (4)$$

, where $lnVA_{it}$ is log value-added and cap_{it} is a set of dummies we create to proxy for capital that we do not observe. In fact, at each survey, firms are asked to rate the state of their technical equipment on a scale from 1 to 5, where 1 is the best. Creating four dummies for each score other than the worst will hence give us a vague indication of a firms capital intensity, which may proxy for the capital stock in a production function.

Table 9: Output Elasticities of Employment

	All	Manufacturing	LBE Manufacturing	Services	LBE Services
starterl	0.689	0.784	0.755	0.620	0.629
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
nonstarterl	0.668	0.754	0.710	0.614	0.609
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
cap1	0.196	0.092	0.045	0.237	0.360
	(0.0004)	(0.2708)	(0.7283)	(0.0338)	(0.2537)
cap2	0.192	0.076	0.040	0.236	0.367
	(0.0005)	(0.3554)	(0.7582)	(0.0347)	(0.2447)
cap3	0.157	0.050	0.017	0.205	0.330
	(0.0043)	(0.5382)	(0.8978)	(0.0654)	(0.294)
cap4	0.057	-0.021	-0.056	0.119	0.227
	(0.2912)	(0.7884)	(0.6534)	(0.2825)	(0.4646)
firm fe			yes		
year fe			yes		
N	122541	24997	13459	74371	33600
R^2	0.085	0.101	0.107	0.072	0.100

Regressions over domestic and switching firms only. P-values in parentheses

The results displayed in table (9) are broadly consistent with our earlier findings. Throughout the subsamples we use for our analysis, we find that the output elasticity of employment is higher once firms have begun to export (the coefficient on variable $starterl$). Looking at sectoral differences, we find an increase of 3 (4.5) percentage points in the manufacturing sector (LBE manufacturing), whereas this increase is 0.6 (2) percentage points in the (LBE) services sector. Unreported results for a regression over switching firms only yield even higher differences in all sectors. The full set of dummies for capital intensity yields significant and economically reasonable results only when regressing over the entire set of firms and in part for the services sector.

5 Conclusion

Our study has revisited the self-selection vs learning-by-exporting debate using detailed data on German firms across all economic sectors. We have exploited variation within and across firms of entering into exporting to gauge whether firms self-select into exporting through higher pre-exporting productivity levels and/or whether firms upgrade their productivity prior to or after entry into exporting. We have also investigated the channels through which productivity effects may occur. We find that future exporters do display higher productivity levels than firms that never export, lending strong support to the self-selection hypothesis. However, average pre-exporting productivity levels remain relatively constant up to entry into exporting, upon which point we register strong increases in productivity. These productivity gains in turn lend strong support to the learning by exporting hypothesis, in that productivity growth picks up only after entry into exporting. This effect is stronger for manufacturing firms than for services firms, in that the former exhibit persistent growth in productivity past entry into exporting, whereas this effect is limited in time (2 years on average) for services firms. We also find that not all sectors display this effect to the same extent. In fact, we have identified a number of subsectors in both manufacturing and services, in which learning by exporting holds, while this effect is not significantly present in others. We explain the different performances of the manufacturing and services sector with significant inherent differences in average propensities to export, which are substantially lower for the services sector. Furthermore, we are able to show that across sectors the size of the LBE effect depends on the level of within-sector competition. In line with basic microeconomic theory, productivity gains are higher for entrants into exporting, which operate in relatively uncompetitive domestic sectors, pointing to an important competitiveness channel for increased productivity through LBE. Moreover, we explain the lower scope for LBE effects in the services sector by uncovering substantially more restrictive access to foreign markets in that sector, which effectively maintains export revenues in only few firms.

Importantly, the overall productivity gains we find are on average not labor-saving, but rather generate increased demand for workers, while basic metrics of working conditions such as the share of temporary and part-time work and average wages do not display particular changes in trend. While we do not investigate policy measures per se, it is safe to conclude from our work that policies aiming at increasing market access may be particularly beneficial for relatively uncompetitive domestic sectors, in terms of productivity gains and employment generation. Notably the services sector displays large asymmetries in available access to foreign markets, which directly translates into lower export-induced productivity gains. While we can make informed statements about the extent of barriers to market access, our data does not allow us to identify their nature. Given the increasing importance of the services sector in generating value-added and employment, further research to highlight what policies contribute to lowering these barriers to foreign market access is of key importance.

References

Aghion, P., Cai, J., Dewatripont, M., Du, L., Harrison, A., and Legros, P. (2015). Industrial policy and competition. *American Economic Journal: Macroeconomics*, 7(4):1–32.

Autor, D., Dorn, D., and Hanson, G. (2013). The china syndrome: Local labor market effects of import competition in the united states. *The American Economic Review*, 103(6):2121–2168.

Autor, D. H. (2003). Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. *Journal of labor economics*, 21(1):1–42.

Bernard, A. B. and Jensen, J. B. (1999). Exceptional exporter performance: cause, effect, or both? *Journal of international economics*, 47(1):1–25.

Bernard, A. B., Jensen, J. B., Redding, S. J., and Schott, P. K. (2012). The empirics of firm heterogeneity and international trade. *Annu. Rev. Econ*, 4:283–313.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics*, 119(1).

Breinlich, H. and Criscuolo, C. (2011). International trade in services: A portrait of importers and exporters. *Journal of International Economics*, 84(2):188–206.

Cernat, L. and Kutlina-Dimitrova, Z. (2014). Thinking in a box: A mode 5 approach to service trade. *Journal of World Trade*, 48(6):1109–1126.

De Loecker, J. (2007). Do exports generate higher productivity? evidence from slovenia. *Journal of international economics*, 73(1):69–98.

De Loecker, J. (2013). Detecting learning by exporting. *American Economic Journal: Microeconomics*, 5(3):1–21.

Eberle, J., Jacobebbinghaus, P., Ludsteck, J., and Witter, J. (2011). Generation of time-consistent industry codes in the face of classification changes: Simple heuristic based on the establishment history panel (bhp). Technical report, Institut für Arbeitsmarkt-und Berufsforschung (IAB), Nürnberg [Institute for Employment Research, Nuremberg, Germany].

Ellguth, P., Kohaut, S., and Möller, I. (2014). The iab establishment panel - methodological essentials and data quality. *Journal for Labour Market Research*, 47(1-2):27–41.

FDZ (2015). Afid-panel strukturerhebung im dienstleistungsbereich. Technical report, FDZ der Statistischen Ämter des Bundes und der Länder.

Fischer, G., Janik, F., Müller, D., and Schmucker, A. (2009). The iab establishment panel-things users should know. *Schmollers Jahrbuch*, 129(1):133–148.

Greenaway, D. and Kneller, R. (2007). Firm heterogeneity, exporting and foreign direct investment*. *The Economic Journal*, 117(517):F134–F161.

Hosono, K., Miyakawa, D., and Takizawa, M. (2015). Learning by export: Does the presence of foreign affiliate companies matter? Technical report, RIETI Discussion Paper, 15-E-053.

Jensen, J. B. et al. (2011). *Global trade in services: fear, facts, and offshoring*. Peterson Institute for International Economics Washington, DC.

Kelle, M. and Kleinert, J. (2010). German firms in service trade. Technical report, Economics working paper/Christian-Albrechts-Universität Kiel, Department of Economics.

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

Lodefalk, M. (2013). Servicification of manufacturing–evidence from sweden. *International Journal of Economics and Business Research*, 6(1):87–113.

Malchow-Møller, N., Munch, J. R., and Skaksen, J. R. (2015). Services trade, goods trade and productivity growth: evidence from a population of private sector firms. *Review of World Economics*, 151(2):197–229.

Manjón, M., Mañez, J. A., Rochina-Barrachina, M. E., and Sanchis-Llopis, J. A. (2013). Reconsidering learning by exporting. *Review of World Economics*, 149(1):5–22.

Melitz, M. J. (2003). The Impact of Trade on Intra-industry Reallocations and Aggre- gate Industry Productivity. *Econometrica*, 71:1695–1725.

Müller, S. (2010). Capital stock approximation with the perpetual inventory method: stata code for the iab establishment panel. fdz methodenreport. *Institute for Employment Research (IAB), Nuremberg*.

Pakes, A. and Olley, S. (1995). A limit theorem for a smooth class of semiparametric estimators. *Journal of Econometrics*, 65(1):295–332.

Stock, J. H. and Watson, M. W. (2008). Heteroskedasticity-robust standard errors for fixed effects panel data regression. *Econometrica*, 76(1):155–174.

Temouri, Y., Vogel, A., and Wagner, J. (2013). Self-selection into export markets by business services firms–evidence from france, germany and the united kingdom. *Structural Change and Economic Dynamics*, 25:146–158.

Van Beveren, I. (2012). Total factor productivity estimation: A practical review. *Journal of economic surveys*, 26(1):98–128.

Van Biesebroeck, J. (2005). Exporting raises productivity in sub-saharan african manufacturing firms. *Journal of International economics*, 67(2):373–391.

Vogel, A. (2011). Exporter performance in the german business services sector. *The Service Industries Journal*, 31(7):1015–1031.

Vogel, A. and Wagner, J. (2011). Robust estimates of exporter productivity premia in german business services enterprises. *Economic and Business Review*, 13(1-2):7–26.

Wagner, J. (2007). Exports and productivity: A survey of the evidence from firm-level data. *The World Economy*, 30(1):60–82.

Wagner, J. (2012). International trade and firm performance: a survey of empirical studies since 2006. *Review of World Economics*, 148(2):235–267.