The Effects of Exporting on Labor Productivity: Evidence from German firms

Johannes Schwarzer *

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 support for the learning-by-exporting hypothesis in manufacturing but less for the services sector, but strong evidence in favor of the hypothesis when considering continuing exporters across both sectors. 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.

*Council on Economic Policies. Email: `js@cepweb.org`

I would like to acknowledge comments and suggestions from Andrea Ariu, Andrew Bernard, Simon Evenett, as well as participants of the second Empirical Investigations in Services Trade workshop at the European University Institute in Florence on earlier drafts of the paper. All remaining errors are my own.

Contents

1 Introduction

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 into and exit from 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. The self-selection hypothesis hence implies that firms which become exporters are simply more productive to start with. There is a broad consensus in the empirical literature as reviewed in Wagner (2007), Greenaway and Kneller (2007a), and Bernard et al. (2012) to this effect, confirming 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. Its early formulations can be traced back to endogenous growth models, such as Grossman and Helpman (1993), who point to technology diffusion through participation in international markets, which may enhance within-firm productivity. Demand side driven exploitation of economies of scale was stressed as an important productivity-enhancing factor in traditional export-led growth hypotheses (eg Kaldor (1970)). Much of the early arguments as made by Pack (1992) or Westphal (2002), however, were based on case studies of the rapidly industrializing Asian Tiger countries. As such, LBE hypotheses often emphasize a variety of mechanisms such as learning from foreign markets in terms of buyer-seller relationships, and increased competition with foreign suppliers, or adapting and improving product quality to suit foreign preferences. While empirical studies generally remain agnostic about the exact mechanism underlying LBE, identification of the effect rests on the assumption that the productivity effect of a firm's international activity must - by definition - be specific to entering foreign markets, entailing activities and knowledge that non-exporters do not possess. The evidence for this effect so far is rather sparse, some examples include Hosono et al. (2015), Fernandes and Isgut (2015), Manjón et al. (2013), Lileeva and Trefler (2010), Bigsten and Gebreeyesus (2009), De Loecker (2007), and Van Biesebroeck (2005), Girma et al. (2004).

A number of recent theoretical developments are consistent with certain aspects of alleged mechanics of the latter effect. Emerging literature on multi-product firms points to adjustments in product mixes as a result of increased competition in export markets, which induces firms to focus on their core competencies and adjust their product offer accordingly, resulting in firm-level productivity gains (eg. Mayer et al. (2014), or **?**). Bustos (2011) models technology adoption jointly with entry in to exporting and finds empirical evidence for trade liberalization induced innovation in both new and existing exporters from Argentina. Importantly, this mechanism seems to hold also in advanced economies: Lileeva and Trefler (2010) find evidence for similar predictions for Canadian exporters following tariff reductions in the US. Verhoogen (2008) associates the decision to export with a joint decision of product quality upgrading to serve consumer preferences in the foreign market. All these papers underline a specific aspect that may help explaining the observed correlation between exporting and productivity. However, the theoretical basis for motivating either effect is often problematic, as the models of heterogeneous firms are generally static. As such, the specific predictions they make are immediately associated with entry into exporting and complicate empirical work that seeks to disentangle pre- and post exporting firm performance. As noted by Tybout (2003), identification of the immediate link between productivity at the time of entry into exporting is often problematic, as the econometrician usually does not have all the necessary information, especially around the time of entry into exporting. In other words, the decision to enter exporting may matter more than the actual entry into exporting. There are two main strategies that have been adopted to remedy this issue. First, depending on the availability of reasonable instruments, the treatment variable (entry into exporting) can be instrumented by a third variable that correlates with export, but not with productivity. Such an approach has been used with a broad variety of context-specific instruments, for example in Van Biesebroeck (2005) (Ethnicity of firm owner and state ownership, Lileeva and Trefler (2010) (tariff cuts), Verhoogen (2008) (Peso devaluation), and Bustos (2011) (lagged tariffs). A second approach that has been used to identify the effect of exporting on productivity is motivated by the fact that the counterfactual productivity trajectory of exporters - *had they not started exporting* - cannot be observed, even though that counterfactual should ideally be the benchmark against which one wants to test potential productivity effects. During the latter 2000s, the trade literature started borrowing from techniques developed in the field of labour economics to construct a proxy for this counterfactual, essentially following the approach developed by Heckman et al. (1997). The approach consists of creating control groups using matching techniques based on observable characteristics. This is the approach followed by Girma et al. (2004) Greenaway and Kneller (2007b), De Loecker (2007) and Bigsten and Gebreeyesus (2009), for example. Even though matching on observables does not solve the potential bias introduced by omitting variables that are unobserved to the econometrician, the approach may help reduce the bias under the assumption that firms that tend to be similar in observables should also be similar in unobservables.

Other recent contributions allow for more explicit post-entry "learning" effects in dynamic setting. Eaton et al. (2009) develop a model of firm-level exporting behavior that takes account of search and learning processes in foreign markets by allowing learning from those markets and its competitors there to identify potential buyers. Freund and Pierola (2010) introduce uncertainty and sunk cost associated with the development of export-market specific products to explain survival patterns of new entrants. Aw et al. (2011) model endogenous R&D decisions jointly with exporting, which can explain post-entry productivity growth of exporters. Albornoz et al. (2012) develop a model which assumes uncertainty about firms' general ability to earn profits abroad, which can be resolved only through trial-and-error experience in foreign markets. Similarly, Timoshenko (2015) develops a model in which learning about foreign demand accounts for product switching in foreign markets. **?** develops a model that stresses the importance of buyer-seller networks and derives predictions about history-dependent export expansion induced by the reduction of search frictions through participation in foreign markets.

Conceding important data limitations as contextualized in section 2, our study uses a robust set of specifications to investigate the LBE effect in Germany, using data on 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.

Our study contributes 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, across all economic sectors. To this end, we are drawing on the methodology developed by Autor (2003), employing it for the first time to analyze the link between exporting and productivity. Our analysis also relates to studies that focus on the interplay between sectoral competition and productivity, such as Greenaway and Kneller (2007b), Mayer et al. (2014), and more generally Aghion et al. (2015). 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 also find increasing productivity rates for continuing exporters. 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.

Section 2 describes the dataset we use in detail and discusses advantages and shortcomings in the context of recent advances in the literature. It also contains discussions of methodological issues and the construction of variables. Section 3 contains a step-by-step approach towards our main analysis in 3.5. We then turn to three different robustness checks in section 4, that address several shortcomings of our analysis from a different methodological angle. We conclude in section 5.

2 Data

We use confidential, representative German establishment-level survey 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)).[2] 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, 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.

The dataset has been extensively used for German labor market research,

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

[2]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).

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.[3]

Advantages of using the IAB dataset in the LBE context

Our dataset offers a unique opportunity to examine firm responses to export participation for firms across all sectors in Germany. The examination of this particular dataset has two major advantages:

First, as outlined in the literature review above, we can jointly examine firm responses across the manufacturing and the services sector, which has not been done before for Germany. Those studies that do exist on Germany have relied on sectoral data, small samples and short time spans. In contrast, our dataset is representative of the entire German economy and covers the entire period after German reunification. In comparison with other datasets on German firms, the IAB panel stands out as the most comprehensive one. 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 that dataset is perhaps the most accurate of all, it only 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 that data with either datasets from the IAB or the German Federal Statistics Office.

The IAB panel offers a consistent set of key outcome variables over time that we can observe across a large panel of firms, without methodological break that could compromise comparability. The dataset merges survey data with administrative records on the classification of economic activity and in-

[3] 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.

cludes a number of further measures to ensure high data quality for scientific purposes, as described later in this section.

Second, we can add to a better understanding of the under-researched case of Germany, and in doing so contribute to the literature on LBE in advanced countries more generally. What makes the analysis of LBE in the German context most interesting is that the German economy is traditionally export-intensive and technologically advanced. As pointed out above in the literature review, the early studies on LBE have been largely devoted to the analysis of developing countries, both conceptually and empirically. Nevertheless, the question of whether access to foreign markets improves firms' performance is equally important for developed economies and therefore highly relevant. The underlying hypothesis has often been one of "technological catch-up" through access to foreign markets, which would seem to apply less to developed economies such as Germany. However, there are at least two reasons why potential LBE effects should not be confined to developing countries only:

(i) Substantial variation of firm characteristics within countries and even industries have been the *raison d'être* for the emergence of the heterogeneous firms literature. Modern trade literature acknowledges that the departure from representative firms models has been a crucial element for the better understanding of economic realities. The literature on multi-product firms takes this a step further and allows for variation in productivity even across products within the same firm. This stylized pattern suggests unequal access to technology across firms regardless of the aggregate economy's stage of industrialization and has led researchers to investigate individual firm responses to trade shocks. As such, the traditional argument for LBE effects can be transposed to the firm level and be equally applied to industrialized economies such as Germany. Indeed, evidence for such export-induced productivity gains and greater technology adoption has been found in the case of Canadian exporters into the US (Lileeva and Trefler (2010)).

(ii) There are a variety of mechanisms underlying the LBE hypothesis, as reviewed in section (1), which are not necessarily specific to developing country exporters. For instance, Mayer et al. (2014) stress product adjustment and competitiveness channels through which exporting may result in productivity gains and find evidence for this channel in French exporting firms. Some dynamic "learning" models reviewed in section (1) identify mechanisms that relate exporters' performances to uncertainty about foreign markets and adjustments to foreign preferences. While their predictions are generally tested against data from developing country exporters, there is no *a priori* reason why they should not hold for developed economies such as Germany, as long as export markets differ from the home market. A Colombian exporter to the US arguably faces greater adjustment than a German exporter, yet the latter will still be learning to survive and thrive in a foreign market as well. Ultimately, the question is an empirical one and comparing domestic producers with exporters of the same country should give us a clearer idea of the existence and the magnitude of these effects. Our dataset hence allows us to investigate this issue and contribute to the discussion of whether LBE effects can be detected in developed economies or not.

Caveats of using the IAB dataset in the LBE context

The advantages of our dataset must be discussed in the light of certain caveats that working with German data entails.

Our analysis will not focus on any particular mechanism that may underpin observed LBE effects. Rather, we aim at establishing a general causal link between entry into exporting and productivity, regardless of product-destination characteristics. This is simply because we cannot distinguish between a firm's products and their characteristics, as well as its export destinations. More generally, we cannot distinguish between services and goods trade, such as eg 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.

We do not have information on imports and hence cannot account for potential import-induced productivity effects that can result from increased import competition or greater availability of intermediate inputs (see eg. Goldberg et al. (2009), Amiti and Khandelwal (2013) and Halpern et al. (2015) for evidence on these channels). We are less concerned about this omission for two main reasons: 1) we do not exploit a specific change in the trade environment to identify the export decision and its productivity effect and 2) we rescale the time variable as described in 2.2, which would disperse any punctual change in the trade environment across different points in the time dimension we generate specifically for each firm we observe.

Lacking information on a firm's capital stock, we cannot estimate traditional measures of total factor productivity (TFP). We discuss this issue and how we deal with it in greater length in section 2.4. At any rate, De Loecker (2011) argues that productivity outcomes need to be analyzed together with market power and prices, implying empirical studies need further information on firms' cost structures and markups. Questioning the large empirical literature that has estimated productivity based on proxying output with sales and exploited trade liberalization periods to identify changes in the trade environment, De Loecker (2011) argues that the relationship between measured productivity and trade liberalization may simply occur through the liberalization's impact on prices and demand, implying that the impact on actual productivity cannot be identified. In order to address this bias, De Loecker et al. (2016) estimate a quantity-based production function using data that contain the prices *and quantities* of firms' products over time, highlighting the need to additionally account for the allocation of inputs and their prices into the product mix of multi-product firms. Lacking information on quantities of firm inputs and outputs, we are hence not able to construct productivity measures based on quantities, products and their inputs. As such, our study relates most closely to traditional studies measuring revenue-based proxies of productivity without distinguishing between single- and multi-product firms. We are aware of this caveat, but aim to refine our results with a joint analysis on sectoral competitiveness after estimating our baseline model. We also seek support for the validity of our productivity measure by checking whether new exporters simultaneously increase their domestic sales, analogously to Lileeva and Trefler (2010), in our robustness checks.

An inherent weakness of survey data as we use it is that it relies on reported data and may hence display inaccuracies in comparison with administrative data. Such inaccuracies may persist despite substantial efforts undertaken by the IAB to ensure high quality of the data. For example, more than 70% of interviews are conducted face-to-face, by highly qualified interviewers, interviewing a member of the executive board in 47.8% of all cases and a head of department in 16.7% (Ellguth et al. (2014)).

Another inherent problem of survey data is that firms drop in and out of our sample for reasons we do not 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 each year. The success of these efforts is evident in the large number of observations listed in table (2), 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. Given our topical focus, it is hence important to mention that 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 tend to re-export during the period of observation and 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. Moreover, should cessation of activities be the prime reason for dropping out of the sample, we would expect declining productivity towards the end of each firm's observation period. Arguably, declining productivity is an important factor leading to firm death. In light of the dynamic nature of our methodological setup, this bias would go against the effect we are seeking to uncover, as evidence for LBE would show in an increasing productivity trajectory.

At any rate, the selection of establishments into the sample follows explicit stratification schemes that are based solely on establishment size, sector of economic activity and geographical location (Ellguth et al. (2014)). Therefore, any establishment opting out of the sample will be replaced by another establishment that is highly similar along those three stratification variables only. Since the design of the IAB survey does in no way pay attention to export status when selecting the firms in the sample, we assume that the existence of either potential bias should not be overstated, as whatever bias that may exist would occur in both exporting and non-exporting plants, and hence both in the treatment, as well as in the control group. The absence of a systematic bias between both should therefore not significantly affect our analysis. In addition, in our empirical strategy, we exploit differences across groups and time in within-firm variation, in order to isolate the net effect of exporting, correcting for any potential time invariant biases on the firm level (such as inherent productivity differences). Moreover, we define our main measure of switching into exporting status in a conservative way that rather biases against LBE, as explained in more detail in section 2.3. Finally, in our robustness checks, we employ an alternative technique to define the control group based on observable similarity, ultimately upholding our baseline results.

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 (1).

Table 1: 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

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. Since entry into exporting does not occur for all firms at the same time, we need to establish a common time scale along which we can compare firm performance. 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. On the other hand, a firm that enters the sample with positive export values but subsequently quits the export market will be assigned the interval value of 0 at the moment the switch takes place, i.e. the first year in which it is not an exporter anymore and all subsequent years.[4] For strictly domestically operating firms that never start to export, the value zero is simply assigned to their rounded mean period of observation. For example, if a non-exporting firm is observed for 4 consecutive years, the first year will be assigned the value of -2, the second year will be -1, the third year will be 0 and the fourth year will be 1. As there is no particular or systematic importance to this moment of time, we do not worry about non-random peculiarities that may significantly confound the comparison of entry into exporting against non-exporters.

2.2.1 Controlling for Calendar Years

While this common scale allows us to compare firms along this time dimension, we need to be aware that we may end up comparing firms at different calendar years. For example, if firm A enters into exporting in 1999 and firm B does so in 2013, both firms are assigned time interval 0. A number of issues may make such a comparison unacceptable without proper controls: Changes in the labor market environment, business cycle, trade policy and advancements in technology are just a few factors to be mentioned here. Therefore, we employ calendar year fixed effects in all our regression specifications, so

[4] In fact, we construct a separate time scale for the quitter variable, such that the switcher variable is defined on one scale and the quitter variable is defined on another, analogous scale. For the sake of readability, we only refer to one single time scale in the text.

as to be sure to control for all calendar year-specific factors that the rescaling of the time variable has cluttered. As hinted at in the data description above, the rescaling of the time variable has the added advantage of dispersing potentially confounding changes in the economic environment across time, which makes us more confident in the validity of using a dataset that does not contain information on (product-)export destinations and prices, as well as imported inputs.

2.3 Definition of firm groups

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, namely the first switch into exporting. We hence 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 subsequent export status. The same applies to firms that export in the first year of observation, but then cease to export. This is our baseline characterization of export entrants for two main reasons: First, if learning from export markets were of any relevance, we would expect firms that have exported to display productivity gains regardless of whether they continue to export or not. Second, such a conservative definition of export entrants hence also helps counter a potential pro-LBE bias that could be inherent in our dataset (see discussion in the data description above and the discussion on measuring productivity below).

Nevertheless, we want to be able to also consider firms that quit exporting. Intuitively, we expect productivity gains to be stronger for continuing exporters once they are purged of quitters. Analogously to the construction of our $STARTER$ variable, we hence construct a variable $QUITTER$ that take on the value of 1 the first time a firm is observed to cease exporting, regardless of whether it resumes exporting at a later time or not. One can think of this variable as the inverse of the $STARTER$ variable.

Altogether, we define three broad groups of firms, defined as:

1. Domestic firms if $\sum_t EXP_{it} = 0$
2. Switchers if $0 < \sum_t EXP_{it} < T$
3. International firms if $\sum_t EXP_{it} = T$

The group of $Switchers$ contains firms for which the $STARTER$ variable takes the value of 1 at least once during their period of observation, but also firms for which the value of $QUITTER$ take the value of 1 at least once. The two subgroups are hence not mutually exclusive. Nevertheless, given our definition, each firm switches only once, whether into starting and/or quitting. Further, if we consider only those firms that are switchers, but who never display a positive value of $QUITTER$, we can focus on those new exporters that - once entering into exporting - remain exporting until the end of their observations, and which we will call $SURVIVORS$.

Given that not each firm is not observed in every time period, we end up with an unbalanced panel whose distribution is concentrated around time 0 and thins out towards the tails. What are the implications of this concentration? Given the calculation of the $STARTER$ and $QUITTER$ variable and the rescaled time variable, we are left with a time span reaching from -21 to +21. Obviously, the further we go back in time ($t < 0$) or ahead of time ($t > 0$), we need to be careful about the judgments we make on the representativity of firm characteristics, as extreme values become more likely due to the lower number of sampled firms, which is why most of our analyses focus on a greatly reduced time span. Additionally, in order to minimize potential biases that accrue from this imbalance, we perform a propensity score matching technique on each time period to identify an average treatment effect in 4.3.

2.4 How to measure productivity?

The standard approach to proxy for firm productivity is to retrieve the residual of a production function, which then in turn is compared across exporting and non-exporting firms. Resulting productivity measures may be biased for four main reasons, of which the first two have a long tradition of debate in the literature (Van Beveren (2012)): i) input choices may be endogenous, a function of firm efficiency, 2) a selection bias due to the exit of firms and input choices made conditional on survival, iii) an omitted price bias when revenue-based productivity estimation is used and prices not properly accounted for (De Loecker (2011)), and iv) a bias resulting from assuming identical production techniques and final demand across products manufactured by a single firm (Bernard et al. (2009)). In the LBE context,De Loecker (2013) shows that estimating production functions that do not allow exporting status to affect productivity - as widely used in the literature on LBE - will eventually translate into a bias against the LBE hypothesis.

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 the discussion for this dataset in Müller (2010)). Instead, we use the different measure of labor productivity, constructing our variable as value-added per worker. Having value-added, rather than output, as the numerator of our labor productivity measure refines the indicator in that an important factor of production - intermediate inputs - is accounted for, at least in value terms. However, a serious drawback remains the lack of an adequate capital measure, in whose absence we cannot account or the contribution of capital to increased per worker value-added. Fixed effects estimation and controlling for investment in all our main specifications cannot fully remedy this shortcoming. We further take up the issue in the robustness checks towards the end of the study, but already caution that any generalization from labour towards TFP based on this study should be taken with a grain of salt. 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.

We can address some of the known issues in measuring productivity in the following ways: When estimating productivity in our main specification (3),

Table 2: Data description

Variable	Obs	Mean	Std. Dev.	Min.	Max.
exp	213931	.2298171	.4207161	0	1
starter	213931	.2923559	.454846	0	1
quitter	213931	.090403	.2867580	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

we use a fixed effects estimator as in Pavcnik (2002) and Petrin and Levinsohn (2003). As such, we can to some extent account for unobserved factors such as time-invariant productivity shocks that may in part be a reflection of unobserved capital stocks that differ across firms. Such an approach also helps to attenuate the simultaneity bias resulting from endogenous input choices Ackerberg et al. (2007) and reduce selection bias resulting from endogenous firm exit in the sample, to the extent that it is related to the time-invariant productivity component. Furthermore, we explicitly take a possible correlation between exporting and productivity into account by including exporting status directly in equation (3), similarly to Bigsten and Gebreeyesus (2009) or Van Biesebroeck (2005).[5]

Concluding the data section, we list the summary statistics of the main variables used in table (2). Our baseline analyses will be based on 156189 observations, as per the productivity variable, which is defined only for firm-year data points for which data on both value-added and employment is available. Exporting status is observed in almost 23% of all cases, which increases to 29.2% for the starter variable, which remains 1 despite subsequent exit of previous exporters. The quitter variable is relatively low, indicating relatively little exit from exporting.

[5]We use current values of the $STARTER$ or $QUITTER$ variables rather than lagged exporting. Section (3.5) will deal more specifically with dynamic variation.

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 start by following Bernard and Jensen (1999) and others in estimating variants of the following equation by Ordinary Least Squares:

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

$prod_{ikt}$ 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 (1). Given the log specification 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. Columns (1) are over the entire sample of firms, columns (2) purge international firms, columns (3) are only over switchers. 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 subsample 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.

In a second exploratory step, we split our group of switchers into its two components, survivors and quitters. We perform analogous regressions and display the results in table (4). The coefficients obtained on surviving firms (those who keep exporting after their first entry) are all larger than those obtained for the entire switcher group in (3), which makes intuitive sense. Likewise, the coefficients of quitters are all smaller, however still significantly large and positive, implying higher average post-quitting productivity even over the restricted sample in columns (3).

Table 4: Pooled OLS: Switcher subgroups

	Survivors						Quitters					
	1a	1b	2a	2b	3a	3b	1a	1b	2a	2b	3a	3b
starter	0.588	0.333	0.479	0.343	0.223	0.148						
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)						
quitter							0.206	0.168	0.323	0.248	0.073	0.046
							(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
labor	NO	YES	NO	YES	NO	YES	NO	YES	NO	YES	NO	YES
N	130616	130616	106466	106466	7143	7143	156189	156189	132039	132039	25573	25573
R^2	0.172	0.195	0.128	0.149	0.107	0.141	0.120	0.169	0.119	0.148	0.060	0.096

Columns (a) refer to regressions without the employment control variable, (b) columns refer to those with the employment variable. Columns (1) are over the entire sample of firms, columns (2) purge international firms, columns (3) are only over the specific subgroup. p-vales in parentheses

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 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. However, the post-entry trajectory is not increasing, which we would expect if LBE were present. Recall the identification problem we have at $t = 0$. In that context, let us also recall the quite restrictive characterization of our $STARTER$ variable, which groups all switchers together, regardless of

Figure 1: Comparing Means: Labor Productivity

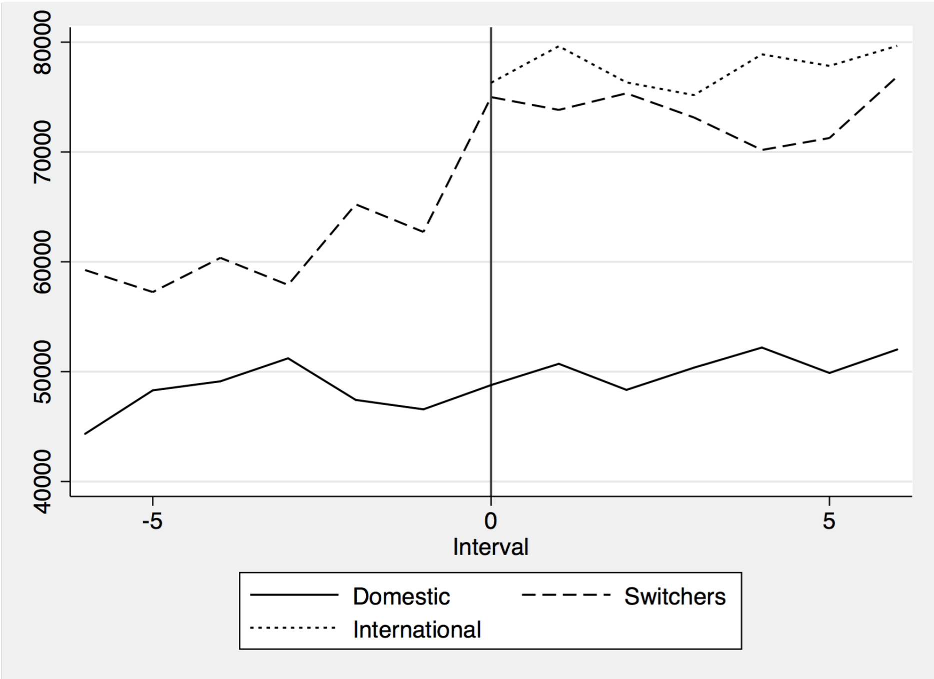

whether they keep exporting or not. If we disentangle this potentially quite heterogeneous group, the trajectories look much more like what we would expect (figure (2)): Continuing exporters have a monotonically increasing post-entry productivity trajectory (except for the last period), while that of quitters seems to be driving the relatively constant post-export entry slope of the aggregate group observed in figure (1). Notice also that quitters experience continuous export growth prior to quitting, i.e. while exporting.[6]

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 (3) and (4).[7] 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,

[6]The time scales for survivors and quitters are not the same. The $STARTER$ variable is defined on one time scale, the $QUITTER$ variable on another, since these characterizations are not mutually exclusive and a firm may be both a starter and a quitter.

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

Figure 2: Comparing Means: Labor Productivity by Subgroup

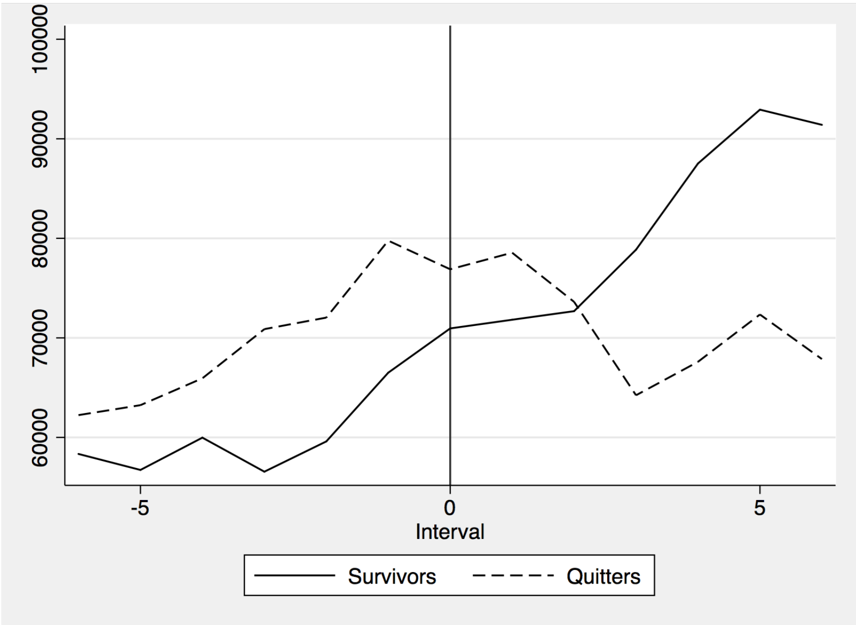

Figure 3: Comparing Means: Employment

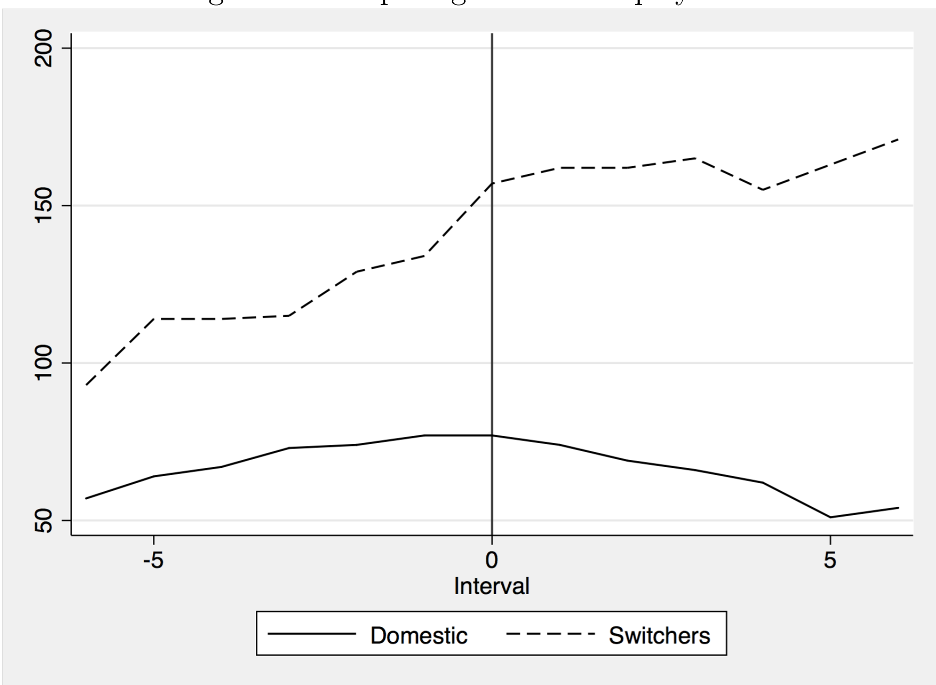

the labor productivity increases that occur with exporting appear to be even

Figure 4: Comparing Means: Inputs per worker

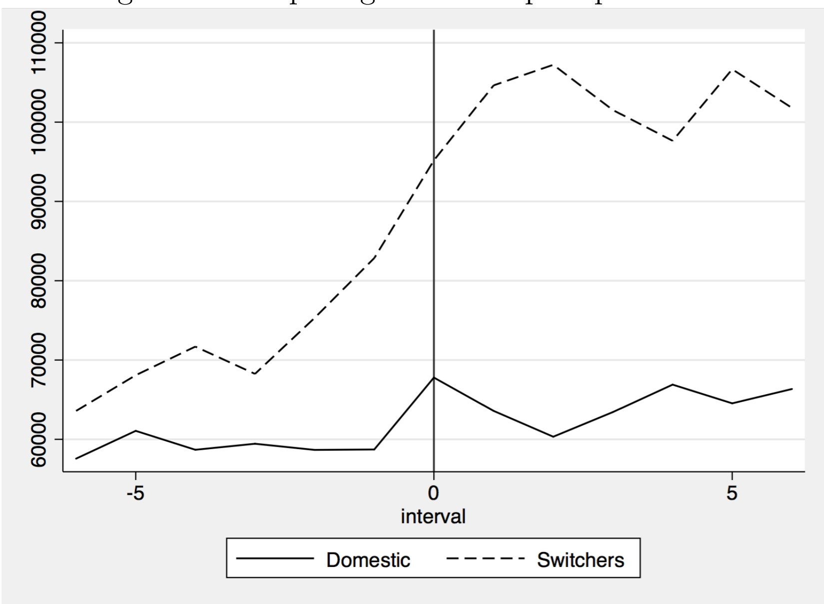

more spectacular and not to be driven by reducing the average workforce. The stark average increase in purchased external inputs confirms larger demand for those per worker. Unreported figures for investment per worker draw a similar picture. Both metrics, however, do not display a substantial pre-exporting 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).[8]

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

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

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

Table 5: 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

Columns (1) and (2) are over the entire sample of firms, column (3) purges international firms, columns (4) is only over switchers. p-values in parentheses.

look at the results in table (5) 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 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

Motivated by our discussion in (2.4), 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 control for firm size - or use of factor inputs - 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, gains that may accrue simply through domestic market expansion.

At this point, it is important to recall that the RHS variables enter the equation solely for the purpose of controlling for time-varying firm-specific characteristics that may matter in terms of productivity. The previous analyses have shown the need to take account of firm specific characteristics, such as in table (3), where we alternate variants of equation (1) with and without inclusion of the employment variable. In the original productivity sense, we are interested in the extent to which firms are able to change their ability to transform inputs into output, or value added in our case. Conditioning on current employment is hence necessary to control for variation in value added that is immediately caused by variation in labor, as the focus of this study lies elsewhere, namely in uncovering the role of exporting. Hence, the coefficient β can be interpreted as a productivity shifter of exporting, conditional on the other controls. With this background, it is crucial to note that having the employment variable on both sides of the equations (1) and (3) does not affect the statistical validity of our analysis. To see this, let us first expand our estimating equation (3) as follows:

$$\begin{aligned}\ln\left(\frac{VA_{it}}{L_{it}}\right) &= \alpha + \beta STARTER_{it} \\ &+ \gamma_1 \ln\left(L_{it}\right) + \gamma_2 \ln\left(\frac{I_{it}}{L_{it}}\right) + \gamma_3 \ln\left(\frac{W_{it}}{L_{it}}\right) + \gamma_4 \ln\left(\frac{D_{it}}{L_{it}}\right) \\ &+ \pi_{year} + \lambda_i + \varepsilon_{it}\end{aligned}$$

Now consider the following equation:

$$\begin{aligned}\ln VA_{it} &= \alpha + \beta STARTER_{it} \\ &+ \delta \ln\left(L_{it}\right) + \gamma_2 \ln\left(I_{it}\right) + \gamma_3 \ln\left(W_{it}\right) + \gamma_4 \ln\left(D_{it}\right) \\ &+ \pi_{year} + \lambda_i + \varepsilon_{it}\end{aligned}$$

The labor variable L_{it} now enters solely the right hand side of the equation. Nevertheless, the two equations are econometrically equivalent, if we let $\delta \equiv (\gamma_1 - \gamma_2 - \gamma_3 - \gamma_4 + 1)$.

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 (6).

Table 6: Fixed Effects Estimation: Sectors

	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

Columns (a) are over the entire sample of firms, columns (b) purges international firms, columns (c) is only over switchers. 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.

Table 7: Fixed Effects Estimation: Subgroups

	Survivors			Quitters		
	a	b	c	a	b	c
starter	0.188	0.250	0.084			
	(0.0000)	0.0000)	(0.0040)			
quitter				-0.042	-0.056	-0.019
				(0.0019)	(0.0000)	(0.2377)
firm controls	yes					
firm fe	yes					
year fe	yes					
N	118987	97564	6435	142448	121025	23461
R^2	0.169	0.268	0.042	0.157	0.207	0.128

Columns (a) are over the entire sample of firms, columns (b) purges international firms, columns (c) is only over the subgroup in question. Standard errors are clustered on the firm level and p-values are given in parentheses

Analysis of figure (2) suggests taking account of different productivity trajectories between surviving exporters and quitters within the switcher group. Varying the control groups analogously to the results in table (6), we perform the analysis separately for survivors and quitters in table (7). The coefficients we obtain for survivors are largely consistent with previous results on switchers. The results for quitters are more interesting. In particular, having added firm level covariates and exploiting within-firm dynamics, we see a reversal of the signs of the coefficients when compared to the specification in table (4). However, this effect is only significant in specifications (a) and (b), where the control group comprises the more productive international and surviving firms. Column (c) attests no notable post-exporting productivity differences within the group of quitters only.

3.4.1 Sectoral Decomposition and Market Structure

Given these results, we 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 (1). 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 (8).

Table 8: 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.

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 concentra-

Table 9: 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

tion, Jensen et al. (2011) obtains a ranking of these sectors according to their 'tradability'. While our sectors listed in table(8) 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 (9) 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 (6). At the same time, table (9) 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 (9), 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.

Figure 5: LBE and Market Concentration

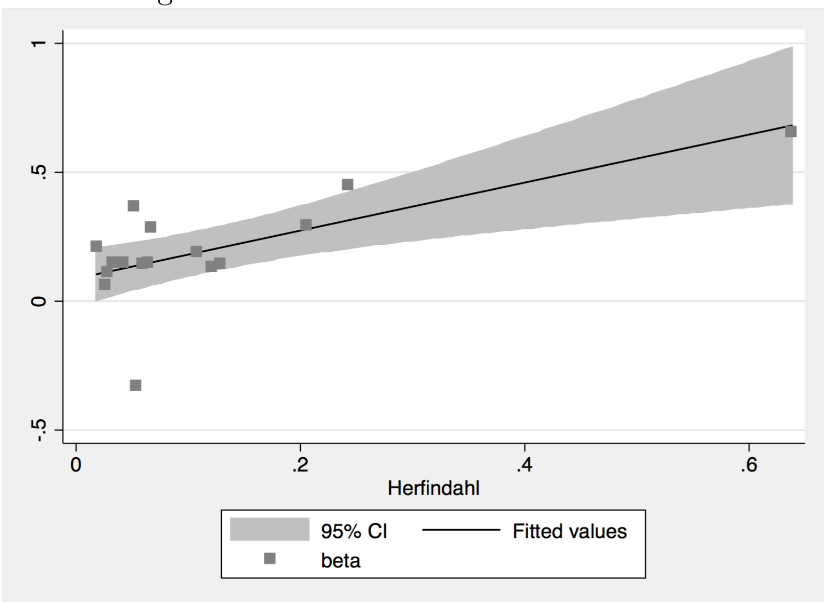

In figure (5) we plot the Herfindahl index we obtain for each sector against the (statistically significant)[9] coefficients we obtain from the indi-

[9] p-value < 0.05

vidual regression results obtained in table(8). 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 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 (9) 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 also observe substantial differences within the group of switchers. As such, firms that continue exporting after their first entry display a greater average post-entry productivity gains that firms that quit exporting.

We have so far attested self-selection to the extent that future exporters are on average more productive than their domestic counterparts. We have also firmly established the result that average post-entry productivity of switchers is higher than average pre-entry productivity. [10]. However, we cannot ascertain yet whether our results are driven by the potentially endogenous year of entry into exporting (see discussion in section (1)).

In order to better account for each firm's productivity trajectory, we follow the method developed by Autor (2003) and augment equation (3) with leads and lags of the $Starter_{it}$ and $Quitter_{it}$ variables, replacing those variables with a set of dummies as follows: We add a dummy for t_{-k}, where k denotes the intervals a firm is observed before entry into (exit from) exporting, as well as a dummy for t_0 and t_{+j}, where j denotes the intervals a firm is observed after exporting (exit from exporting). These dummies each take the value of one only for the year of their corresponding time period and are zero otherwise, which allows us to isolate the average effect in each time period that is being considered. 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

[10]This result is strengthened for continuing exporters, as seen in table (4)

$k = 1$ and $j = 1$ in table (10), 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 criterion of at least one observation prior to exporting (exit from exporting), as well as at least two observations after the year of exporting (exit from 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 (8) 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 and increasing coefficients on $t_{\geq 0}$.

Table 10: Leads and Lags: Switchers

	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 (10) 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 $k = 2$.[11] 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 section (3.4.1).

Table 11: Leads and Lags: Survivors

	All	Manufacturing	LBE Manufacturing	Services	LBE Services
t_{-1}	-0.120	-0.053	-0.088	-0.147	-0.115
	(0.026)	(0.442)	(0.448)	(0.065)	(0.338)
t_0	0.132	0.158	0.305	0.087	0.286
	(0.009)	(0.009)	(0.001)	(0.359)	(0.029)
t_{+1}	0.208	0.246	0.308	0.178	0.305
	(0.000)	(0.000)	(0.000)	(0.029)	(0.020)
$t_{+(2)}$	0.248	0.315	0.404	0.104	0.341
	(0.000)	(0.000)	(0.000)	(0.225)	(0.002)
firm controls			yes		
firm fe			yes		
year fe			yes		
N	88691	14917	7901	56567	26091
R^2	0.273	0.197	0.184	0.266	0.334

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

We now decompose the analysis by considering the two subgroups of switching firms separately. We display the results of the same regression with leads and lags for survivors in table (11) and for quitters in table (12). Consider first table (11). Except for the general services sector, all coefficients reflect LBE to the letter in that each post-entry period is associated with an

[11]Results are not reported

Table 12: Leads and Lags: Quitters

	All	Manufacturing	LBE Manufacturing	Services	LBE Services
t_{-1}	0.053	0.036	0.080	0.062	0.123
	(0.002)	(0.190)	(0.006)	(0.018)	(0.001)
t_0	-0.040	0.018	0.059	-0.062	-0.027
	(0.008)	(0.524)	(0.130)	(0.020)	(0.473)
t_{+1}	0.011	0.062	0.114	-0.027	0.021
	(0.560)	(0.034)	(0.002)	(0.314)	(0.584)
$t_{+(2)}$	-0.016	0.038	0.087	-0.055	0.011
	(0.383)	(0.211)	(0.013)	(0.049)	(0.792)
firm controls			yes		
firm fe			yes		
year fe			yes		
N	130224	36419	19918	71040	32579
R^2	0.159	0.105	0.106	0.167	0.200

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

increasing gain in productivity. Even if endogeneity in the entry year were present, we can confirm productivity gains in each subsequent year. In order to strengthen this result, we compare the lead and lag coefficients in table (13) and report the corresponding p-value of rejecting the null hypothesis of equality of coefficients. The first row tests for equality of the t_{-1} coefficient with the contemporaneous coefficient t_0. In subsequent rows, we test with respect to the other two post-export (post exit from export) coefficients. In the case of both Switchers and Survivors, we can reject the null hypothesis of pairwise equality of coefficients. Looking at quitters in table (12), the exit from exporting does not seem to be strongly related with systematic productivity effects across subsectors. The last period of exporting t_{-1} does indicate higher productivity, but the periods after exit do not display a consistent and significant pattern, notably across sectors. It is important to keep in mind that, analogously to the group of switchers, the group of quitters also comprises firms that start exporting again at a later point in time. By definition, however, the *Quitter* variable remains 1 once a firm exits from exporting.

4 Robustness Checks

4.1 Marginal Labor Productivity

Our analysis has consistently focused on average productivity effects. Here we estimate a simple production function to check for differential marginal

Table 13: Testing equality of coefficients

	Switchers			Survivors			Quitters		
	All	Man.	Ser.	All	Man.	Ser.	All	Man.	Ser.
$t_{-1} - t_0 = 0$	0.0000	0.0000	0.0000	0.0000	0.0021	0.0124	0.0000	0.5187	0.0000
$t_{-1} - t_{+1} = 0$	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0020	0.3448	0.0014
$t_{-1} - t_{+(2)} = 0$	0.0000	0.0000	0.0748	0.0000	0.0000	0.0100	0.0002	0.9407	0.0000

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:

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

, 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.

The results displayed in table (14) 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 statistically significant and economically reasonable results only when regressing over the entire set of firms and in part for the services sector.

Table 14: 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

4.2 Mark-ups and post-export domestic expansion

As discussed above, our approach is inherently prone to productivity mismeasurement, notably due to the lack of data on capital stock. One way to check for the implications of this omission is to check for post-export domestic expansion, as proposed by Lileeva and Trefler (2010). The reasoning is as follows: If we were to pick up higher measured productivity through higher prices fetched in foreign markets and hence charging higher mark-ups, then - absent underlying differences in TFP performance - this would cause exporters to lose customers domestically. In this context, De Loecker and Warzynski (2012) find evidence for increasing mark-ups with export entry of Slovenian firms that may hence be mistaken for productivity improvements, if not properly accounted for. However, within exporters, they do not find significant differences between mark-ups charged domestically and abroad. This is important for the validity of our robustness check, because it suggests that the intuition developed by Lileeva and Trefler (2010) holds, since firms do not apply different mark-ups at home and abroad.

Table (15) displays the result across the now familiar control groups, of a regression of log domestic sales on the *Starter* and *Quitter* variable respectively. Our switchers (and notably survivors) have increased their domestic market sales relative to non-exporters, as well as relative to their pre-export

Table 15: Domestic Sales Expansion

	Switchers			Quitters			Survivors		
	a	b	c	a	b	c	a	b	c
starter	1.456	1.112	0.457				1.897	1.417	0.434
	(0.0000)	(0.0000)	(0.0000)				(0.0000)	(0.0000)	(0.0000)
quitter				0.586	0.959	0.416			
				(0.0000)	(0.0000)	(0.0000)			
industry fe	yes								
year fe	yes								
N	180877	153791	37376	180877	153791	29152	151725	124639	8224
R^2	0.212	0.170	0.120	0.152	0.152	0.126	0.234	0.160	0.125

(a) All firms, (b) No international firms, (c) only subgroup under consideration (Switchers, Quitters, Survivors). Standard errors are clustered on the firm level and p-values are given in parentheses

levels. This result is inconsistent with rising mark-ups, but consistent with higher TFP. It suggests that switchers have indeed gained in TFP and were therefore able to outperform their domestic competitors. The rise in productivity we observe is also reflected in a rise in domestic sales across all subsamples, which strengthens our confidence in the validity of our estimations and our productivity measures.

4.3 Propensity Score Matching

Finally, our last robustness check addresses the possibility of having chosen inadequate control groups, addressing the selection problem in our analysis from a different angle. Recall that in most of our specifications, we compare the effect of export entry (exit from exporting) relative to non-exporters. However, this comparison may not be valid if non-exporters display fundamentally different key characteristics. By employing firm-level fixed effects, we do account for time-invariant heterogeneity across firms. However, a propensity score matching approach based on Heckman et al. (1997) allows us to match exporters with non-exporters on observable characteristics, which helps us proxy for the counterfactual of what an exporter's characteristics would be *had it not started exporting.* The average treatment effect we seek to uncover is given by:

$$\begin{aligned} E[prod_{it}^1 - prod_{it}^0 | Exporter_{it} = 1] = \; & E[prod_{it}^1 | Exporter_{it} = 1] \\ & - E[prod_{it}^0 | Exporter_{it} = 1] \end{aligned} \tag{5}$$

where $prod_{it}^0$ and $prod_{it}^1$ stand for our measure of labor productivity before and after entry into (exporting from) exporting respectively. The variable

$Exporter_{it}$ equals 1 for all switchers both before and after their entry into, or exit from exporting, so as to allow comparisons at intervals $t < 0$, where our standard *Starter* and *Quitter* variables would take on the value of zero. We do not consider international firms for this analysis. The unobservable last term, the counterfactual for each firm, is hence proxied with $E[prod^0_{it}|Exporter_{it} = 0]$, where the corresponding control firm is matched based on the nearest neighbor in terms of its propensity score. We estimate the latter *for each time interval*, so as to account for the changing composition of firms that is the result of having an unbalanced panel and our rescaled time variable as per 2.2. Our matching variables are the same that we use as controls throughout our regressions, but we match within each industry as per 2.1 only, so as to maximize comparability in terms of heterogeneous industry-level export potential.[12] The latter is warranted by the diverging propensities of exporting we observe empirically across industries. The average treatment effects for each group of firms in each interval is displayed in table (16).

Table 16: Average Treatment Effects

	-2	-1	0	1	2
Switchers	-0.002	0.001	0.247	0.190	0.180
	(0.970)	(0.970)	(0.000)	(0.000)	(0.000)
Survivors	-0.084	-0.031	0.110	0.346	0.408
	(0.403)	(0.464)	(0.013)	(0.000)	(0.000)
Quitters	0.148	0.227	0.009	0.102	0.073
	(0.000)	(0.000)	(0.808)	(0.000)	(0.011)

P-values are given in parentheses

Altogether, the results from matching corroborate our previous results in 3.5 quite neatly. Pre-entry differences are insignificant for both Switchers and Survivors, whereas they become significant at the year of entry. For Survivors, the difference is increasing over time. For Quitters, the differences are significant and increasing prior to exit ($t < 0$), but become inconsistent thereafter.

[12]Sectors "Households" and "Extra-territorial Organizations" cannot be considered due to insufficient numbers of observations

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. In contrast, we find strong evidence in favor of the hypothesis when considering continuing exporters across both sectors. 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.

Our results are robust to different specifications and, 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

Ackerberg, D., Benkard, C. L., Berry, S., and Pakes, A. (2007). Econometric tools for analyzing market outcomes. *Handbook of econometrics*, 6:4171–4276.

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.

Albornoz, F., Pardo, H. F. C., Corcos, G., and Ornelas, E. (2012). Sequential exporting. *Journal of International Economics*, 88(1):17–31.

Amiti, M. and Khandelwal, A. K. (2013). Import competition and quality upgrading. *Review of Economics and Statistics*, 95(2):476–490.

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.

Aw, B. Y., Roberts, M. J., and Xu, D. Y. (2011). Rd investment, exporting, and productivity dynamics. *American Economic Review*, 101(4):1312–44.

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.

Bernard, A. B., Redding, S. J., and Schott, P. K. (2009). Products and productivity. *The Scandinavian Journal of Economics*, 111(4):681–709.

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

Bigsten, A. and Gebreeyesus, M. (2009). Firm productivity and exports: Evidence from ethiopian manufacturing. *The Journal of Development Studies*, 45(10):1594–1614.

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

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

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

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

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

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

Eaton, J., Eslava, M., Krizan, C. J., Kugler, M., and Tybout, J. (2009). A search and learning model of export dynamics. *Penn State University mimeo.*

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.

Fernandes, A. M. and Isgut, A. E. (2015). Learning-by-exporting effects: Are they for real? *Emerging Markets Finance and Trade*, 51(1):65–89.

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

Freund, C. and Pierola, M. D. (2010). Export entrepreneurs: evidence from peru. *Policy Research Working Paper Series*.

Girma, S., Greenaway, D., and Kneller, R. (2004). Does exporting increase productivity? a microeconometric analysis of matched firms. *Review of International Economics*, 12(5):855–866.

Goldberg, P., Khandelwal, A., Pavcnik, N., and Topalova, P. (2009). Trade liberalization and new imported inputs. *American Economic Review*, 99(2):494–500.

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

Greenaway, D. and Kneller, R. (2007b). Industry differences in the effect of export market entry: learning by exporting? *Review of World Economics*, 143(3):416–432.

Grossman, G. M. and Helpman, E. (1993). *Innovation and growth in the global economy*. MIT press.

Halpern, L., Koren, M., and Szeidl, A. (2015). Imported inputs and productivity. *American Economic Review*, 105(12):3660–3703.

Heckman, J. J., Ichimura, H., and Todd, P. E. (1997). Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *The review of economic studies*, 64(4):605–654.

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.

Kaldor, N. (1970). The case for regional policies. *Scottish journal of political economy*, 17(3):337–348.

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 productivity… 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.

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

Melitz, M. J. (2003). The Impact of Trade on Intra-industry 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.*

Pack, H. (1992). New perspectives on industrial growth in taiwan. *Taiwan: From developing to mature economy*, pages 73–120.

Pavcnik, N. (2002). Trade liberalization, exit, and productivity improvements: Evidence from chilean plants. *The Review of Economic Studies*, 69(238):245.

Petrin, A. K. and Levinsohn, J. A. (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies*, 70(2):317–341.

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.

Timoshenko, O. A. (2015). Product switching in a model of learning. *Journal of International Economics*, 95(2):233–249.

Tybout, J. R. (2003). Plant-and firm-level evidence on “new” trade theories. *Handbook of international trade*, 1:388–415.

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.

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

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.

Westphal, L. E. (2002). Technology strategies for economic development in a fast changing global economy. *Economics of innovation and new Technology*, 11(4-5):275–320.