BANK FOR INTERNATIONAL SETTLEMENTS

BIS Working Papers
No 469

When firms and industries matter: understanding the sources of productivity growth

by Ulf Lewrick, Lukas Mohler, Rolf Weder

Monetary and Economic Department

October 2014

JEL classification: F14, L60, O47

Keywords: Growth, total factor productivity, inter-industry trade, intra-industry trade

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

When firms and industries matter: understanding the sources of productivity growth[☆]

Ulf Lewrick[b,1], Lukas Mohler[a,2], Rolf Weder[a,3]

[a] *Faculty of Business and Economics, University of Basel, Peter Merian-Weg 6, CH-4002 Basel, Switzerland*
[b] *Bank for International Settlements, Centralbahnplatz 2, CH-4002 Basel, Switzerland*

Abstract

This paper presents a framework to assess the relative importance of three key sources of productivity growth that research on international trade focuses on: (i) inter-industry specialisation; (ii) intra-industry reallocation of resources across heterogeneous firms, including firm entry and exit; and (iii) technological progress. Detailed data on Swiss manufacturing firms illustrate how the framework can be empirically applied. Based on this example, we find that intra-industry reallocations are the most important source of growth in aggregate total factor productivity, reflecting in particular the productivity growth of large, incumbent firms and the entry of new firms. That said, inter-industry specialisation and general technological progress remain important supplementary sources of growth in Swiss manufacturing.

Keywords: Growth, total factor productivity, inter-industry trade, intra-industry trade
JEL: F14, L60, O47

[☆] The views expressed in this article are those of the authors and do not necessarily reflect those of the Bank for International Settlements. Financial support from the Swiss National Science Foundation (Project No. 100014-124975) is gratefully acknowledged, and we are thankful to the Swiss Federal Statistical Office for granting access to the data. For their helpful comments, we also thank Ben Tomlin, Ingo Fender and seminar participants at the Bank for International Settlements as well as at the annual conferences of the Canadian Economics Association, the European Trade Study Group and the Verein für Socialpolitik. Rolf Weder thanks Simon Fraser University and the University of British Columbia for the kind hospitality when working on the paper.

Email addresses: `ulf.lewrick@bis.org` (Ulf Lewrick), `lukas.mohler@unibas.ch` (Lukas Mohler), `rolf.weder@unibas.ch` (Rolf Weder)

[1] Phone: +41 61 280 94 58.

[2] Phone: +41 61 267 07 70.

[3] Phone: +41 61 267 33 55.

1. Introduction

Productivity growth is a driving force for economic welfare. As measures to enhance an economy's productivity rank high on any economic policy agenda, numerous studies have focused on the sources of productivity growth. International trade research is no exception, with different strands of the literature emphasising the role of distinct contributory factors. In broad terms, classical or neoclassical trade models underscore the role of specialisation *across industries* in fostering aggregate productivity, while the so-called new (or "new new") trade theory highlights the impact of reallocations of resources among firms *within industries.* Yet another strand of the literature concentrates on the impact of innovations due to R&D and the role of international spillovers.[4]

How important are these different factors in quantitative terms? To shed some light on this issue, we develop a framework that decomposes the contributions to aggregate productivity growth from (i) an inter-industry effect; (ii) an intra-industry effect; and (iii) general technological progress. This lets us assess whether growth has been predominantly driven by a country's specialisation on its most productive industries, akin to the (neo-)classical trade theories, or whether reallocations within industries and within firms have had an important impact as emphasised by the new trade theory including its extensions with firm heterogeneity.

The distinction is also relevant from a policy perspective. Reallocations across industries, for example, tend to be associated with stronger distributional effects than those within an industry.[5] This may call for a different set of policy responses to mitigate the impact on labour markets. Furthermore, productivity growth that is mainly based on an improved allocation of resources (eg due to the exit of unproductive firms), but which lacks more broad-based advances in technology, may prove unsustainable in the longer term.

An analysis of the Swiss manufacturing sector serves to illustrate how to apply our framework. This sector experienced a strong increase in international trade from 1997 to 2007, suggesting a prominent role for trade-related productivity drivers. We find that intra-industry

[4] Since it is beyond the scope of this paper to provide an overview of the existing literature, we limit ourselves to a small and incomplete selection. Recent contributions in the trade literature that emphasise the importance of the classical or neoclassical trade theories include Eaton and Kortum (2002), Debaere (2003), Costinot and Donaldson (2012) and Finicelli et al. (2013). Important contributions to the new trade theory, including, for example Krugman (1979), Krugman (1980), Melitz (2003) and Bernard et al. (2003), are discussed in Melitz and Trefler (2012). Research on R&D and international spillovers has been sparked by the theoretical contributions of Romer (1990), Grossman and Helpman (1991), and Aghion and Howitt (1992) and related empirical analysis conducted by, for example, Coe et al. (2009). Keller (2004) provides an overview of this literature. Finally, a growing number of studies combine elements of both the traditional and the new trade theories as, for example, in Bernard et al. (2007), Demidova (2008) and Bombardini et al. (2012), or link elements of the new trade theory to technological progress (see Schröder and Sørensen (2012) for a recent example.)

[5] Several studies show that switching industries involves higher costs for workers than staying within the same industry (see, eg, Shin (1997) and Greenaway et al. (2000) for studies of the U.S. and U.K. labour market, respectively). Similarly, from an international trade perspective, trade liberalisation among economies that trade similar products (intra-industry trade) is associated with a weaker impact on factor markets (often referred to as the "smooth adjustment hypothesis") than between economies where liberalisation induces specialisation towards industries of comparative advantage (see Brülhart (2000) for a discussion and empirical analysis).

reallocations are a key factor in explaining changes in Swiss manufacturing productivity. Reallocations within industries account for 54% of the increase in aggregate productivity from 1997 to 2007 based on total factor productivity (TFP) estimates. Furthermore, these reallocations are the most import determinant of year-on-year productivity changes, explaining 64% of the variation in aggregate TFP. Productivity growth within large firms and from firm entry represents the key determinant of the intra-industry effect, although data limitations call for some caution in interpreting the different components of intra-industry contributions. Inter-industry reallocations play a notable role as well. They account for 11% of the overall increase in Swiss manufacturing TFP in the 1997–2007 period. Interestingly, inter-industry reallocations are found to be most pronounced during recessions, pointing to sizeable structural adjustments during such economic downturns as the 2008–09 global crisis. Finally, general technological improvement, which we approximate by the change in the unweighted average of the continuing firms' productivity, accounts for the remaining 35% of growth and explains about one third of year-on-year variation in aggregate TFP.

Our framework builds on several elements of previous research. Our decomposition of aggregate productivity builds on the approaches developed in Baily et al. (1992), Grilliches and Regev (1995), Foster et al. (2001) and Foster et al. (2006). These studies contributed significantly to clarifying the implications for aggregate productivity growth of firm entry and exit as well as of changes in existing firms' market share and productivity. We make two innovations with a view to mapping the different growth channels identified in the international trade literature into distinct components of aggregate productivity growth.

First, we extend the static concept of Olley and Pakes (1996) – that decomposes productivity of one sector – to a dynamic multi-industry dimension. Taking into account multiple industries in a dynamic setting allows us to evaluate the relative strength of inter-industry as opposed to intra-industry adjustments. The advantage of the Olley and Pakes (1996) concept lies in the distinction between changes in industry means and those that are associated with changes in the distribution of firm-level productivity and market shares, a distinction which is closely related to the notion of firm heterogeneity in the new trade theory. Relying on cross-sectional differences in productivity, as noted in Foster et al. (2001), should also mitigate any impact from measurement error and transitory shocks that may impede the accuracy of measures tracking individual firms over time.

Second, we refine the concept of distinguishing between the productivity contributions of different types of firm. We compare entering and exiting firms' contributions to aggregate productivity with those of continuing firms rather than, as in other studies, with that of an aggregate average. This improves the accuracy of measuring exit and entry effects. In this regard, our decomposition is also related to the study of Slovenian firms by Melitz and Polanec (2012), who propose a similar decomposition of the approach in Olley and Pakes (1996). Unlike the specification proposed by Melitz and Polanec (2012), however, the multi-industry setup in our paper allows us to identify reallocations across industries and benchmark productivity changes of the different types of firm to developments in the corresponding industry. Given the sizeable differences in productivity developments across industries, this multi-industry dimension provides important insights for our empirical application.

In the empirical analysis, we combine elements of the methods proposed in Olley and Pakes (1996), Klette and Griliches (1996) and Ackerberg et al. (2006). This enables us to address a number of potential biases looming in conventional firm-level TFP estimates.

The rest of the paper consists of four sections. The analytical framework to decompose aggregate productivity into different sources of productivity growth is presented in Section 2 of this paper. In Section 3, we describe the data and how we estimate firm-level TFP. Section 4 discusses the results of the decomposition of productivity growth in the Swiss manufacturing sector, including a discussion of alternative specifications. Section 5 concludes. The Appendix provides further details of the derivation of the analytical framework and the TFP estimation.

2. The sources of productivity growth

This section derives and discusses the building blocks of the productivity decomposition. To start, we define aggregate productivity at the manufacturing level, Φ_t, as the weighted average of the productivities of all manufacturing firms.

$$\Phi_t = \sum_{j=1}^{J} \sum_{i=1}^{N_{jt}} s_{ijt}\, \varphi_{ijt}, \tag{1}$$

where we sum over all active firms N_{jt} over all industries j at time t. The term φ_{ijt} represents a firm's productivity and s_{ijt} the firm's market share in total manufacturing (eg measured in terms of sales). The shares of all firms in an industry sum up to $S_{jt} = \sum_i s_{ijt}$, and industry shares, S_{jt}, accordingly sum up to unity. In the remainder of the paper, we use the abbreviated notation $\sum_i$ and $\sum_j$ to facilitate the exposition.

The key to assessing changes in aggregate productivity is to recognise that aggregate productivity in equation (1) can be expressed in terms of weighted industry means and deviations thereof as shown by Olley and Pakes (1996, p.1290). Our extension to the multi-industry dimension yields $\Phi_t = \sum_j \sum_i S_{jt}\, (\bar{s}_{jt} + \Delta s_{ijt})\, (\bar{\varphi}_{jt} + \Delta\varphi_{ijt})$. The variables $\bar{s}_{jt}$ and $\bar{\varphi}_{ijt}$ represent the unweighted mean market share $(1\,/\,N_{jt})$ and industry j's unweighted mean productivity, $(\sum_i \varphi_{ijt})\,/\,N_{jt}$, respectively. Δs_{ijt} measures the difference between each firm's market share within its industry and the unweighted mean sales share in the same industry. By analogy, $\Delta\varphi_{ijt}$ equals the deviation of the firm's productivity from the unweighted mean of the corresponding industry. Multiplying out and eliminating terms that are equal to 0 yields the following relationship:

$$\Phi_t = \sum_j S_{jt}\, \bar{\varphi}_{jt} + \sum_j \sum_i S_{jt}\, \Delta s_{ijt} \Delta\varphi_{ijt}, \tag{2}$$

where $\Delta s_{ijt} = (s_{ijt}/\, S_{jt}) - \bar{s}_{jt}$ and $\Delta\varphi_{ijt} = \varphi_{ijt} - \bar{\varphi}_{jt}$.

Aggregate productivity is equal to the sum of two elements. First, a measure of the overall productivity level. This is measured by the weighted sum of the average firm productivity of all industries (the first term on the right-hand side). Second, a measure of the allocative efficiency of the resources used in manufacturing, as captured by the sum of the covariances between firms' market share and industry productivities (the second term). This term yields a positive sign whenever a firm with (below-) above-average productivity accounts for (a below-) an above-average market share as compared with the industry's mean.

2.1. Towards a trade-related decomposition

We now turn to the decomposition of *changes* in aggregate productivity. To align the decomposition with the productivity drivers discussed in the international trade literature, we disentangle the impact of inter-industry reallocations from those occurring within industries. As a first step, we denote the change in aggregate productivity from period $t = 0$ to t as (see Appendix A for more details):

$$\begin{aligned} \Phi_t - \Phi_0 &= \sum_j (S_{jt} - S_{j0}) \big[(\sum_i \Delta s_{ijt} \Delta\varphi_{ijt}) + \bar{\varphi}_{jt} \big] \\ &+ \sum_j S_{j0} (\bar{\varphi}_{jt} - \bar{\varphi}_{j0}) + \sum_j S_{j0} \big(\sum_i \Delta s_{ijt} \Delta\varphi_{ijt} - \sum_i \Delta s_{ij0} \Delta\varphi_{ij0} \big). \end{aligned} \tag{3}$$

The first term on the right-hand side measures changes in aggregate productivity due to changes in inter-industry specialisation over time, using current firms' market share and productivities. Note that the term in squared brackets is equal to the industry equivalent of total manufacturing productivity as can be seen by inspecting equation (2).[6] The second term, $\sum_j S_{j0} (\bar{\varphi}_{jt} - \bar{\varphi}_{j0})$, captures changes in average productivity. The third term focuses on changes in firms' market share and productivity within industries. For both the second and the third term, we hold industry weights constant at their levels in $t = 0$, since changes in these weights are accounted for by our measure of inter-industry specialisation (ie the first term in equation (3)).

As a second step, we distinguish between continuing firms (indicated by capital letter C), exiting firms (X) and entering firms (E). This enables us to quantify the impact of firm entry and exit as well as reallocations among continuing firms within industries that represent the main transmission channels of welfare gains according to the new trade theory literature. In addition, we are able to assess the impact of technology changes abstracting from changes in the composition of active firms.

For each category of firms – C, E and X – we define distinct productivity averages at the industry level as $\bar{\varphi}^{\gamma}_{jt} = (\sum_{i\in\gamma} \varphi_{ijt}) / N^{\gamma}_{jt}$ for $\gamma = \{C, E, X\}$, so that, for example, $\bar{\varphi}^{C}_{jt}$ represents

[6] To interpret reallocations across industries relative to a common benchmark, a helpful adjustment is to rewrite $\sum_j (S_{jt} - S_{j0}) [(\sum_i \Delta s_{ijt} \Delta\varphi_{ijt}) + \bar{\varphi}_{jt}]$ as $\sum_j (S_{jt} - S_{j0}) \Delta \Phi_{jt}$. Here, we define $\Delta \Phi_{jt}$ as the deviation of industry j's productivity, Φ_{jt}, from the unweighted mean of industries' productivity levels $\bar{\Phi}_{jt} = J^{-1} \sum_j \Phi_{jt}$ (for further details see Appendix A).

the unweighted average of all continuing firms' productivity of sector j at time t. Then, based on equation (3), we decompose the change in aggregate productivity into an *inter-industry* effect, a *technology* effect and an *intra-industry* effect. The intra-industry effect itself is an aggregate of four effects that capture the industry dynamics of the three types of firms – continuing, entering and exiting ones:

$$\begin{aligned}
\Phi_t - \Phi_0 \;=\;& \underbrace{\sum_j (S_{jt} - S_{j0})\,\Delta\Phi_{jt}}_{\text{inter-industry effect}} + \underbrace{\sum_j S_{j0}\left(\bar{\varphi}^C_{jt} - \bar{\varphi}^C_{j0}\right)}_{\text{technology effect}} \\
&+ \underbrace{\sum_j \sum_{i\in C} S_{j0}\,\Delta\, s_{ijt}\,(\Delta\,\varphi_{ijt} - \Delta\,\varphi_{ij0})}_{\text{within-firm productivity effect}} + \underbrace{\sum_j \sum_{i\in C} S_{j0}\,\Delta\,\varphi_{ij0}\,(\Delta\, s_{ijt} - \Delta\, s_{ij0})}_{\text{within-firm share effect}} \\
&+ \underbrace{\sum_j S_{j0}\, s^E_{jt}\,\left(\bar{\varphi}^E_{jt} - \bar{\varphi}^C_{jt}\right) + \sum_j \sum_{i\in E} S_{j0}\,\Delta s_{ijt}\Delta\varphi_{ijt}}_{\text{firm-entry effect}} \\
&- \underbrace{\sum_j S_{j0}\, s^X_{j0}\,\left(\bar{\varphi}^X_{j0} - \bar{\varphi}^C_{j0}\right) - \sum_j \sum_{i\in X} S_{j0}\,\Delta s_{ij0}\Delta\varphi_{ij0}}_{\text{firm-exit effect}}\,. \qquad (4)
\end{aligned}$$

The *inter-industry effect* has two elements. The first one, $S_{jt} - S_{j0}$, is a measure of the share of industry sales that has been shifted to or away from an industry since $t = 0$. The second one, $\Delta\,\Phi_{jt}$, indicates whether this shift has been to an industry with above- or below-average productivity as compared with the industries' average. The inter-industry effect will thus turn positive, whenever the country's manufacturing sector has, on average, shifted resources towards more productive industries.[7]

The *technology effect* captures the overall trend in industry-level productivities, abstracting from firm size and excluding the impact of market entry and exit by including continuing firms only. It serves as an indicator of broad-based technological improvements affecting all firms in an industry. Changes in the manufacturing sector's mean productivity are controlled for shifts across industries by weighting averages at each industry's original market share, S_{j0}.

The *intra-industry effect* is composed of the remaining *four effects* highlighted in equation (4). The *within-firm productivity effect* accounts for the improvement in allocative efficiency resulting from relatively large (small) firms, in terms of market share, raising their productivity by more (less) than the average firm in the industry. Note that this effect takes up differences in the technological development among firms as emphasised by the new trade theory. The *within-firm share effect* measures changes in allocative efficiency that are due to a reallocation of market shares at the firm level for a given productivity level. This effect empirically replicates the reallocation gains suggested by standard heterogeneous firm models,

[7] To keep the exposition straightforward, we assume the number of industries, J, remains constant from $t = 0$ to t. This assumption is likely to be in line with most firm-level data sets.

based on market shares being shifted towards the most productive firms. Taken together, these two within-firm effects measure the continuing firms' contribution to allocative efficiency, ie, $\sum_j \sum_{i \in C} S_{j0} \, (\Delta \, s_{ijt} \, \Delta \, \varphi_{ijt} \, - \, \Delta \, s_{ij0} \, \Delta \, \varphi_{ij0})$.

The *firm-entry effect* and *firm-exit effect* assess the impact of market entry and exit on aggregate productivity developments, respectively. Both effects have two components. The first one benchmarks the average productivity of entering (exiting) firms against that of (and only of) continuing firms. This appears to be a more informative benchmark than using (initial or average) aggregate productivity levels as in, for example, Grilliches and Regev (1995) or Foster et al. (2006), as we compare two mutually exclusive sets of firms in the same period rather than a subset of firms to an average or past sample of firms. Melitz and Polanec (2012) show that using this alternative benchmark has a substantial effect on the entry and exit effects proposed in their study of Slovenian manufacturing firms. Unlike their approach, however, our decomposition benchmarks firm productivities at the industry level to account for the significant difference in industry characteristics. By measuring the share of entering (s^E_{jt}) and exiting (s^X_{j0}) firms relative to the total number of active firms of the corresponding period and industry, the component appropriately weights the impact on average productivity resulting from firm entry and exit, respectively. The second component, the covariance term, captures the impact of entering (exiting) firms on the allocative efficiency of the industry: if relatively large and productive firms enter (exit), overall productivity rises (falls). This component complements the assessment of firm entry and exit proposed in earlier studies by taking account of the effect of changes in the composition of active firms.

2.2. Discussion

The proposed decomposition provides a unifying accounting base to quantify and distinguish the different transmission channels emphasised by the main trade theories. The inter-industry effect captures what the traditional trade theory would emphasise, ie aggregate productivity increases due to specialisation towards industries with a comparative advantage. The intra-industry effect not only depicts the adjustment within industries among heterogeneous firms studied by the new trade theory; with its four components, it also allows for differentiating between the channels through which the industry dynamics translate into changes in aggregate productivity. Finally, the technology effect captures efficiency gains of all firms in an industry associated with trade-induced innovation through technology spillovers or competitive pressure as emphasised by the literature on trade and endogenous growth.

The four components of the intra-industry effect are helpful when interpreting recent findings in the new trade literature. *First*, there is evidence for the most productive firms to self-select into export markets and increase their size (see, for exmaple, Bernard et al. (2007)). This would be measured by our within-firm share effect and lead to an improvement in allocative efficiency.

Second, there is inconclusive evidence as to whether these firms also experience productivity gains from entering these markets as discussed by Greenaway and Kneller (2007) and

Wagner (2007). Studying the impact of Mercosur on Argentinean exporters, Bustos (2011), for example, finds support for her model in which the most productive firms (and equivalently the relatively large exporters) upgrade their technology. This would be captured by the within-firm productivity effect.

Note that the within-firm productivity effect may turn out to be negative if, for example, relatively unproductive firms with below-average market shares improved their technology and raised productivity by more than their peers as found, eg by Lileeva and Trefler (2010) when studying the impact of the Canada-U.S. Free Trade Agreement on Canadian manufacturing firms. They describe this effect as "negative selection" which would, in our decomposition, reduce the allocative efficiency measured by the intra-industry effect as resources shift, in principle, to the "wrong" firms. This outcome would, however, be associated with a positive technology effect, as the average productivity of the continuing firms' rises in this instance.

Third, with respect to the firm-entry effect, empirical patterns studied for Spanish (Farinas and Ruano (2005)) and German manufacturers (Wagner (2010)) suggest that entering firms often exhibit lower revenue-based productivity levels than incumbent firms, whereas Foster et al. (2008) find the opposite result for U.S. manufacturing plants with productivity measures based on physical output. In addition to testing these patterns, our framework also allows us to assess the impact on aggregate productivity in terms of the efficiency of resource allocation. Even though the entry of highly productive firms tends to raise aggregate productivity through the first term of the firm-entry effect (see equation (4)), the combined entry effect may nevertheless turn out to be negative. This is the case if these firms were to have below-average market shares such that the allocative efficiency component (the second term of the entry effect in equation (4))) becomes sufficiently negative.

Finally, regarding the firm-exit effect, exiting firms are typically marked by below-average productivity levels in theoretical as well as empirical analyses (see, eg Bartelsman and Doms (2000)). Note, however, that – by analogy to firm entry – the firm-exit effect may reduce aggregate productivity if exiting firms with below-average productivity also have market shares that are sufficiently smaller than the industry average. The overall effect on aggregate productivity, in turn, depends on which firms pick up the market share of the failing firms.

3. Empirical application

This section presents an empirical application of the framework derived above by illustrating the different sources of productivity growth in Swiss manufacturing. We first describe the data and outline the estimation of firm-level TFP, one of our productivity measures, before going on to discuss the results of our analysis in the next section.

3.1. Data

Our data set was provided by the Swiss Federal Statistics Office (SFSO) and consists of an unbalanced panel of 3,564 Swiss manufacturing firms for all years from 1997 to 2009.[8] The data set comprises all manufacturing industries and roughly one quarter of all active firms with more than 10 employees. These firms employ nearly half of the manufacturing sector's workforce and contribute to about 60% of their industry's sales on average.

Each firm is assigned to one of 22 different industries according to the two-digit Swiss industry classification (ie NOGA 2002, exactly corresponding to ISIC 3.1). We group these into 13 industries to combine similar industries and achieve samples of meaningful size for our TFP estimations (see below). Industries include, for example, "Chemical and Pharmaceutical Products", "Machinery and Equipment" or "Textiles, Apparel and Leather Products". Additional data such as aggregate industry statistics and deflators were obtained from publicly available sources, including the SFSO, OECD and the Swiss Federal Customs Administration. While the sales shares of our sample at the industry level are close to the industry shares published by the SFSO including all firms in the industry, they do not perfectly match. To ensure that our aggregate results are representative of the Swiss manufacturing sector, we rescale the sample industry shares to match the shares reported by the SFSO in our calculations. This adjustment reduces primarily the weight of the chemical and pharmaceutical industry for the most recent years in our sample but leaves the main results of our analysis unaffected.

Entry and exit rates in our sample are about 6% and 4% on average, if we exclude the years in which the Swiss economy witnessed a recession (ie 2003 and 2009). These rates are somewhat higher than those that can be derived from the Swiss Business Census and supplementary SFSO data for the years from 2001 to 2008 (which, however, also include the energy industry), ie about 2% (entry) and 3% (exit) on average. Yet, for the recession years, a careful inspection of our data shows that the sample seems to considerably overestimate the sector's entry and exit rates.

Even though the focus of our exercise is to illustrate how to apply the productivity decomposition empirically, rather than highlighting any particular developments in Swiss manufacturing, this feature of the sample is an obvious obstacle to a precise measurement. That said, any resulting bias is likely to primarily affect the individual components of the intra-industry effect with a more limited impact on the overall intra-industry effect. This is because the technology effect is based on the unweighted average of continuing firms' productivity, which should prove relatively robust to changes in the composition of the sample as long as the number of continuing firms remains sufficiently large relative to the number of entering and exiting firms. Since the inter-industry effect is agnostic to the classification of firms, the overall intra-industry effect – being the residual of the two other effects – should prove robust to these measurement issues. Nevertheless, interpreting the results particularly for the recession years requires some caution.

[8] To the best of our knowledge, the data set is unique for Switzerland in terms of coverage of firm-level information, allowing for a comprehensive assessment of productivity developments. Access to the firm-level profit & loss statements and balance sheet data were granted by the SFSO subject to confidentiality requirements.

Finally, as to the definition of the key variables, we use firm-level data on revenue (deflated by industry-level revenue indices from SFSO), employment (full-time equivalents (FTE)), labour productivity (deflated revenue divided by FTE), average wages (deflated labour expenditures including social security divided by FTE, using a country-level wage index), capital (deflated stock of movables and machinery, using import indices for investment goods), intermediates (deflated nominal material expenditures, using OECD intermediate input indices), investment (changes in the deflated capital stock) and industry production (deflated industry-level revenue, using industry-level revenue indices).

Table 1: Summary statistics

Year	**1997**		**2001**		**2005**		**2009**	
Sales[1]	59	(76.8)	67	(75.7)	71	(77.9)	91	(94.3)
Average wage[2]	77	(1.2)	83	(1.2)	89	(1.3)	93	(1.2)
Employment[3]	177	(45.7)	180	(42.7)	171	(38.7)	194	(42.8)
Capital/employee[4]	21	(1.3)	23	(1.6)	23	(2.3)	27	(2.0)
Intermediates/employee[4]	132	(1.7)	143	(1.5)	157	(1.8)	153	(1.5)
Labour productivity[4]	270	(1.5)	296	(1.5)	322	(1.8)	324	(1.6)
Number of firms	1754		1931		1955		1657	

Notes: These are sample averages for selected years. Ratio of the average value of the highest ($q_{80,100}$) to lowest quintile ($q_{0,20}$) of firms ranked by employment in parenthesis. [1]: In millions of Swiss francs. [2]: In thousands of Swiss francs. [3]: In full-time equivalents. [4]: In thousands of Swiss francs per full-time equivalent.

Table 1 shows summary statistics for the firms' main characteristics in 1997 and every fourth year up to 2009. In addition, we compare the sample averages of the highest to lowest quintile of firms ranked by employment. The data reveal a typical pattern that is in line with many other studies focusing on firms in international trade (see, eg, Bernard et al. (2007)): Larger firms pay higher (average) wages, employ more capital and spend more on intermediates per employee, and are more productive in terms of labour productivity (measured as each firm's deflated revenue divided by the number of employees).

On average, the labour productivity (LAP) of the largest firms was 51% above that of the smallest firms in 1997. Notably, the largest firms' productivity advantage substantially increased after 2001, supported by stronger reliance on intermediates and capital per employee than their smaller competitors. This trend coincides with the implementation of the first Swiss-EU Bilateral Agreement, signed in 1999, which was followed by a second agreement signed in 2004. These agreements span a broad set of measures covering, among others, the establishment of labour mobility, the abolition of technical barriers to trade, and transport policy. While the combined impact of these measures is likely to be significant, the complex interlinkages and the fact that firms largely anticipated their introduction prevent a precise quantification. So rather than estimating the impact of trade liberalisation, we interpret our results as predominantly reflecting the increased economic integration of the Swiss manufacturing sector.

The recession of 2009 marks a break in the overall trend, with the productivity advantage of large firms falling back to earlier levels. This suggests that large, internationally active Swiss manufacturers felt the impact of the 2009 recession most strongly, as suggested by the unprecedented decline in both Swiss imports (–10%) and exports (–15%) during this year.[9] Note that the higher and relatively rising LAP of large firms is partly due to the fact that these firms tend to use relatively more of capital and intermediates. Hence, firm-level estimates of TFP that control for observed differences among firms may provide a more accurate view of productivity developments. We thus base our empirical application on the decomposition of changes in TFP, while considering developments in LAP in the discussion of alternative specifications.

3.2. TFP estimates

We start from a definition of productivity at the firm level to avoid any ambiguity. We assume firm i produces a unit of output, Q_{ijt}, at time t in industry $j \in J$, using three factors of production, ie capital (K_{ijt}), labour (L_{ijt}) and intermediates (M_{ijt}). Production is given by a Cobb-Douglas production function, $Q_{ijt} = \varphi_{ijt} \, K_{ijt}^{\alpha_k} \, L_{ijt}^{\alpha_l} \, M_{ijt}^{\alpha_m}$, with the α's representing the elasticity of substitution of the respective input factors. We consider two productivity measures in this paper: first, the firm's TFP which is defined by φ_{ijt} and, second, the firm's LAP which is equal to $Q_{ijt} \, / \, L_{ijt}$.

Van Beveren (2012) provides a review of the potential sources that may bias the coefficient estimates of the production function. The data at hand allow us to address three of them: The bias arising from (i) the endogeneity of the firm's input choices; (ii) sample selection; and (iii) omitting firm-level output prices. Whereas the correction for the bias arising from (i) and (ii) is quite common, only a few studies take into account the bias from (iii).[10]

To do so, we combine three different estimations procedures. First, to account for the endogeneity in the firm's input choices, we base our estimation on the procedure proposed by Ackerberg et al. (2006), henceforth referred to as ACF. This procedure is robust to identification issues potentially arising in the widely used procedures of Olley and Pakes (1996) and Levinsohn and Petrin (2003). Second, we incorporate the control for sample selection as initially developed by Olley and Pakes (1996) into the ACF framework. Finally, as we expect a high level of product differentiation and thus price variation within industries in our data, we combine the above with the strategy proposed by Klette and Griliches (1996) to address the potential bias from omitting firm-specific output prices. Our TFP estimates thus take into account the estimated markup on variable costs in each industry. We refer the reader to Appendix B for a detailed account of the TFP estimation method.[11]

[9]These figures represent the year-on-year decline in real terms. By comparison, Switzerland's real GDP declined by approximately 2% in 2009.

[10]Notable exceptions are, inter alia, Ornaghi (2006) and De Loecker (2011).

[11]A great variety of techniques to estimate TFP are described in the literature. In this paper, we limit ourselves to parametric and semi-parametric estimators and refer to Van Biesebroeck (2007) and Van Biesebroeck (2008) for a discussion of non-parametric techniques such as index numbers and data envelopment analysis.

Columns 1–3 of Table 2 present the coefficient estimates of our benchmark production function (ACF). The first column reports the results without correcting for the omitted price bias. In the second column, we show the results including this correction, assuming a common markup across industries. The third column shows results allowing for industry-specific markups. We correct for sample selection in all three specifications. For comparison, we report the coefficient estimates for the OLS fixed-effects regression (fourth column) and the estimates resulting from the procedures of Olley and Pakes (1996) and Levinsohn and Petrin (2003), referred to as OP and LP.

Overall, the coefficient estimates are in line with the predictions from theory.[12] Comparing the first and the fourth column, we find that by considering the endogeneity of firms' input decision to productivity shocks the ACF specification mitigates the upward (downward) bias in the labour (capital) coefficient in the OLS estimation. That said, the capital coefficient remains relatively low, corroborating the findings of comparable estimates for other small open European economies.[13]

Table 2: Production function estimates

	ACF			**OLS**	**OP**		**LP**	
Labor	0.485	0.499	0.505	0.521	0.514	0.485	0.455	0.450
	(0.006)	(0.006)	(0.008)	(0.024)	(0.006)	(0.006)	(0.012)	(0.018)
Capital	0.068	0.058	0.035	0.007**	0.052	0.075	0.066	0.064
	(0.003)	(0.004)	(0.005)	(0.003)	(0.002)	(0.001)	(0.017)	(0.017)
Intermediates	0.468	0.438	0.481	0.348	0.484	0.443	0.464	0.461
	(0.003)	(0.003)	(0.003)	(0.021)	(0.003)	(0.003)	(0.082)	(0.071)
Industry prod.		0.154	ISM			0.176		0.054
			(0.003)			(0.003)		(0.007)
Obs.	17263	17263	17263	22627	17263	17263	22627	22627

Notes: ** $p < 0.05$; all other coefficients are statistically significant at $p < 0.01$. Standard errors in parenthesis. ACF: Ackerberg et al. (2006); OP: Olley and Pakes (1996); LP: Levinsohn and Petrin (2003); ISM: Regression with industry-specific markups.

Turning from the first to the second column, we note that the coefficient on industry production is statistically significant and has the expected sign which supports our strategy for correcting for the omitted price bias. This is also true for the corresponding specifications based on OP or LP (fifth and seventh column). The implied markups range from a rather timid 6% (LP) to the more reasonable values of 18% (ACF) to 21% (OP) for these specifications.[14]

[12] See, for example, De Loecker (2011) for a more detailed discussion of the theoretical predictions.

[13] See, for examples, Van Beveren (2012), De Loecker (2011) for Belgium, Lööf and Andersson (2010) for Sweden and Görg et al. (2008) for Ireland. As noted in De Loecker (2011), the coefficient on capital measures the elasticity of an input considered to be fixed with respect to the firm's reaction to a contemporaneous productivity shock. We might thus expect the coefficient to be small.

[14] From equation (12) in Appendix B we know that the estimated elasticity of substitution is given by $\hat{\sigma}_j = 1 / \hat{\beta}_j$. The estimated markups are thus equal to $\hat{\sigma}_j / (\hat{\sigma}_j - 1) = 1 / (1 - \hat{\beta}_j)$.

In the third column, we present the results when loosening the assumption of a common markup across industries to capture the heterogeneity across industries. This yields similar headline coefficient estimates (as reported in the table), but introduces larger differences in estimated elasticities across industries when scaling up according to the estimated industry markups.

In the remainder of the paper, we rely on the TFP estimates generated by the ACF algorithm including controls for omitted price bias using industry-specific markups and sample selection (Table 2, third column) as these are likely to account for any identification issues arising from OP or LP. That said, our results prove robust to the selection of any specific two-stage estimation algorithm.

4. Results

We apply the proposed productivity decomposition using the firm-level TFP estimates to assess the contributions to productivity growth from technological progress as well as from reallocations within and across industries. As reallocations across industries may unfold at a slower pace than within industries, we focus on the cumulative growth – rather than yearly changes – of the manufacturing sector's TFP starting from the year 1997. We first present the results of comparing the years 1997 to 2007, before going on to discuss the effects observed during the recessions.

4.1. Technology, inter-industry and intra-industry effects

Figure 1 illustrates the growth path of aggregate TFP in the Swiss manufacturing sector based on the TFP estimates derived from the ACF algorithm and using revenues to calculate firms' market share (s_{ijt}). Based on a decade of observations, TFP rose by 45% from 1997 to 2007, amounting to an annual growth rate of 3.8%. All three components of the productivity decomposition as defined in equation (4) prove to be important sources of aggregate productivity growth. The red bars indicate the contribution from the inter-industry effect, the yellow bars that of the intra-industry effect and the blue bars show the technology effect.

A number of observations stand out. First, while reallocations matter both within and across industries, intra-industry effects tend to contribute more to aggregate productivity changes in the Swiss manufacturing sector than the inter-industry effect. Indeed, 54% (or 24 percentage points) of the cumulative growth in aggregate TFP from 1997 to 2007 is attributable to intra-industry reallocations, whereas inter-industry ones only accounted for 11% (or 5 percentage points). It is only since the 2003 recession that we find reallocations across industries contributing to productivity growth, coinciding with the implementation phase of the first Swiss-EU Bilateral Agreement.

Figure 1: Sources of TFP growth

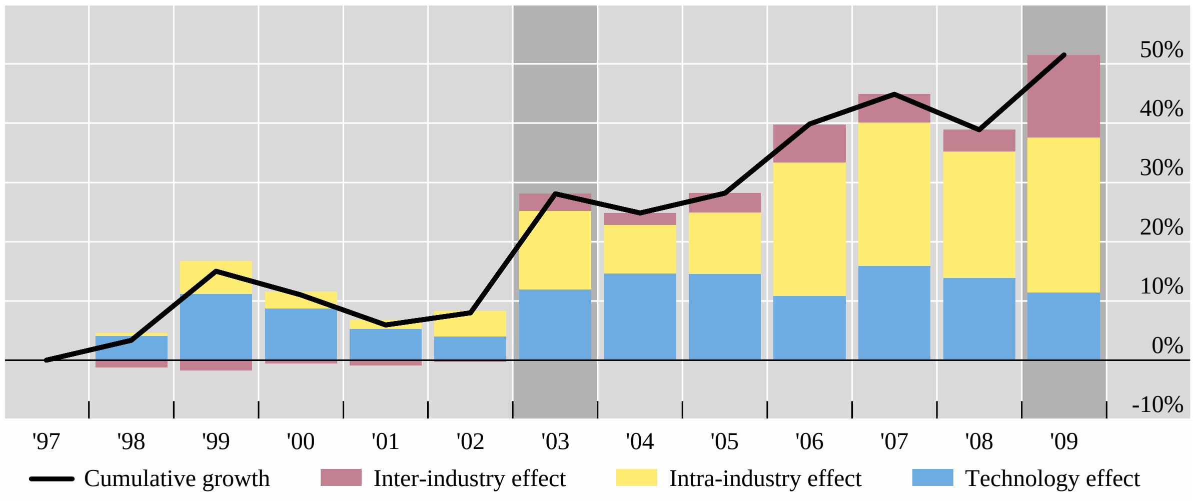

Second, growth in the average productivity of continuing firms, ie the technology effect, assumes an important role. For 2007, the cumulative technology effect contributed 35% (or 16 percentage points) to aggregate productivity growth, although it has been relatively stagnant for the most recent sample years. As we would expect, the pace of continuing firms' technological advancement is notably slower in our sample and thus contributes less to aggregate productivity growth than in sectors and countries experiencing major transitions. For the Slovenian manufacturing sector, for example, the findings in De Loecker and Konings (2006) and Melitz and Polanec (2012) suggest that productivity growth within continuing firms was the predominant source of aggregate TFP growth during the country's transition period from 1995 to 2000.

Third, economic downturns seem to have a strong impact on the relative strength of productivity drivers. Even though we observe a rise in aggregate productivity during both recessions (2003 and 2009, as highlighted by the grey regions), the contributing factors differ. In 2003, the increase in aggregate TFP rests on a positive contribution from all three effects. In 2009, by contrast, the average productivity of continuing firms declines (ie a negative technology effect). This effect is counterbalanced by sizeable intra- and inter-industry reallocations, suggesting considerable structural adjustments towards relatively more productive firms and industries. One important contributing factor to the inter-industry effect, for example, has been the sizeable shift in revenue shares towards the chemical and pharmaceutical industry (which has among the highest TFP levels in Swiss manufacturing) during this period.

4.2. Zooming in on the intra-industry effect

To provide additional insights into the intra-industry dynamics, we illustrate in Figure 2 the contributions from productivity and share effects within firms (blue and yellow bars) as well as those from firm exit and entry (red and green bars) based on equation (4). This decomposition highlights that the net contribution of the intra-industry effect to aggregate productivity growth (as shown in Figure 1) masks opposing trends in its underlying components.

Figure 2: Decomposition of the intra-industry effect

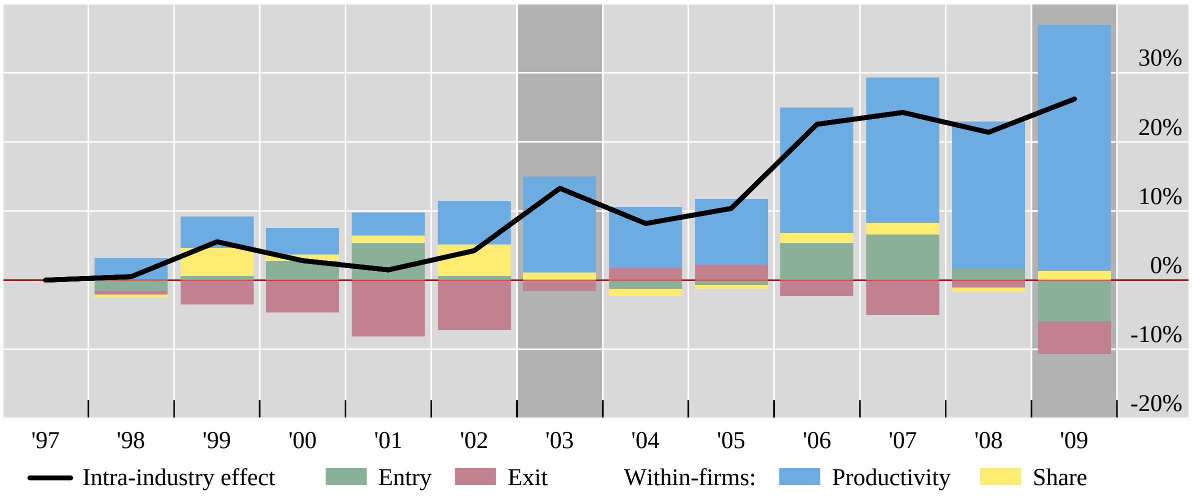

Within industries, improving the allocation of resources among continuing firms turns out to be the most important source of productivity growth in our sample. Indeed, productivity enhancements at relatively large firms (as measured by the within-firm productivity effect) contributed 21 percentage points to the cumulative growth in aggregate TFP from 1997 to 2007, with shifts in market share towards the most productive firms (share effect) adding another 2 percentage points. This result resembles the findings in Baily et al. (1992, p. 207) who report a dominant contribution from continuing plants to aggregate productivity growth in US manufacturing (1972–87).

The entry of new firms contributes nearly 7 percentage points to aggregate TFP growth from 1997 to 2007. Looking into the determinants of the entry effect (see equation (4)) we note that entering firms raised aggregate productivity mainly by improving allocative efficiency. Even though these firms often exhibited below-average productivity, they improved allocative efficiency since they typically accounted for below-average market shares. Exiting firms were similarly characterised by productivity and market share that were both below-average. Their exit thus negatively affected allocative efficiency, with the effect not being fully compensated by an improvement in the average productivity of surviving firms (ie the first term of the firm-exit effect) in most years.

One open question, however, relates to the potential impact of the measurement issues discussed in Section 3.1. Trade theory provides some indications of the expected sign of any potential bias. Since the Swiss manufacturing sector was marked by increasing international integration during the period of observation, our presumption based on heterogeneous firms models (eg Melitz (2003)) is that the least productive firms that account for the smallest revenue shares are the most likely to have exited the market. If firms that no longer report data have not exited the market (and are thus erroneously identified as exiting firms), these would tend to be more productive than firms that have truly exited. In this case, the first term of the exit effect in equation (4) would overestimate the decline in aggregate productivity (or underestimate its growth) that is due to firm exit. The impact on the second term of the exit effect (ie the associated allocative efficiency) depends on the sign of $\Delta\, s_{ij0}\, \Delta\, \varphi_{ij0}$. Even though it is not possible to gauge the sign of this term for the (weighted) sum of all potentially misspecified exiting firms, any resulting (over-) underestimation of the exit effect will be reflected in an equivalent (under-) overestimation of the combined within-firm effect leaving the other components of the framework unaffected.

Similar considerations can help evaluate any potential bias in the entry effect. Indeed, our finding of a rather limited contribution from firm entry to aggregate productivity growth differs from the results found for U.S. manufacturing (see Foster et al. (2001)) or other sectors such as retail trade (see Foster et al. (2006)). Assuming that entrants exhibit below-average revenue-based TFP, as suggested by the sample data but also by findings in, for example, Farinas and Ruano (2005) and Wagner (2010), we would expect continuing firms that are misspecified as entering ones to raise the first term of the entry effect (see equation (4)). One explanation for our finding of relatively limited entry effects could then be due to missing observations on small highly productive firms (for which $\Delta\, s_{ijt}\, \Delta\, \varphi_{ijt} \,<\, 0$) that are accordingly misspecified as entering ones. Yet, this downward bias would be counterbalanced by an upward shift in the combined within-firm effect, preserving the accuracy of the overall intra-industry effect.[15]

4.3. Contributions to the variation of aggregate productivity

We complement our analysis by examining what share of variation in aggregate productivity is explained by the inter-industry, intra-industry and technology effects. The first column of Table 3 presents these shares, based on separately regressing the year-on-year changes of each of the three effects on the corresponding changes in aggregate TFP.

Intra-industry effects are the key driver of (short-term) variation in aggregate productivity for the Swiss manufacturing sector, explaining 64% of the variation in TFP, whereas the inter-industry effect proves statistically insignificant. Yearly changes in aggregate productivity are thus strongly driven by reallocations within industries rather than across industries. This supports the view that reallocations across industries unfold at a slower pace, with limited explanatory power for short-term productivity growth. By comparison, the technology effect contributes 31% to the variation in aggregate TFP. This appears to be less than in the analysis of U.S. manufacturing (1977–96) presented in Petrin et al. (2011), where gains in the technical efficiency of all active firms are the main source of volatility in aggregate productivity.

[15] Another potential issue relates to the measurement of TFP given the absence of information on output and input prices at the firm-level. Foster et al. (2008), for example, show that entrants in selected U.S. manufacturing industries charge exceptionally low prices, resulting in artificially low revenue-based TFP estimates. Such a bias would tend to overstate the contribution to aggregate TFP from continuing firms. By comparison, Mairesse and Jaumandreu (2005) report a very small bias in their study of French and Spanish firms. In contrast to these studies, Atalay (2014) estimates TFP with reference to both firm-level output and materials prices. His study suggests that accounting for materials prices reduces the contribution of net entry to productivity growth, if compared with measures that only account for output prices. While our data do not allow us to account for firm-specific prices, De Loecker (2011) argues that as long as higher input prices are reflected in higher output prices, taking account of the omitted price bias (as done in our study) will at least partially remedy any bias.

Table 3: Contributions to the variation of aggregate productivity (1997–2009)

	TFP		**LAP**	
Weights	sales	employment	sales	employment
Intra-industry effect	0.642***	0.574***	0.841***	0.622**
	(0.084)	(0.085)	(0.079)	(0.158)
Inter-industry effect	0.048	0.046	0.123	0.002
	(0.072)	(0.040)	(0.085)	(0.123)
Technology effect	0.309**	0.379**	0.036	0.376
	(0.097)	(0.107)	(0.132)	(0.240)

Notes: ** $p < 0.05$, *** $p < 0.01$. Standard errors in parenthesis. The results for TFP are based on pooled regressions using ACF with industry-specific markup estimates.

4.4. Alternative specifications and productivity measures

One important question is whether our findings are robust to alternative specifications. As we have shown in Table 2, the different two-stage estimation algorithms yield fairly similar coefficient estimates across a variety of specifications so that selecting a different algorithm than ACF leaves the results qualitatively unaffected.

A more profound robustness check is to change the basis of the firm-level weights (s_{ijt}). Instead of applying output weights (ie sales shares) to aggregate the TFP estimates, we assess the impact of using input weights, such as employment shares. While this reduces our estimate of the cumulative growth in aggregate TFP to 22% (1997–2007), the relative importance of the different effects is largely preserved. Using this alternative specification, inter-industry reallocations account for 9% of aggregate TFP growth, as compared with 11% for our benchmark estimate in Section 4. The intra-industry effect, by comparison, now accounts for 39% of aggregate TFP growth. This confirms our finding of more pronounced intra-industry than inter-industry reallocations in Swiss manufacturing. That said, the relative importance of reallocations declines in this alternative specification if compared with the contribution from growth in the average productivity of continuing firms. The latter, as measured by the technology effect, contributes 52% to aggregate TFP growth if employment shares are used, as compared with 35% in the benchmark estimate. This shift is also reflected in the greater contribution of technology effects to the variation in aggregate TFP (see Table 3, second column), when switching from sales to employment weights.

In the last part of this section, we consider developments in labour productivity (LAP) as an alternative measure of productivity developments. The left-hand panel of Figure 3 illustrates the cumulative growth in aggregate LAP and its decomposition into the three main effects using revenue shares to weight firm-level productivities. The right-hand panel provides further detail on the intra-industry effect by showing the developments in its four different components.

Starting with the decomposition of aggregate LAP growth (left-hand panel), we note that the overall developments broadly resemble those for TFP. Cumulative growth in Swiss manufacturing LAP from 1997 to 2007 stands at 56% with the difference to TFP growth (45%)

reflecting greater use of capital and intermediates per employee as discussed in Section 3.1. The latter also helps explain a more important contribution from the technology effect (53% or 29 percentage points) at the expense of the intra-industry effect (33% or 19 percentage points). While the average LAP of continuing firms rises with greater capital and intermediate use over time (as captured by the technology effect), the impact of this trend is largely offset in the intra-industry effect where the underlying components compare firm-level LAPs of the same year. Finally, the relative contribution to aggregate growth from inter-industry effects remains broadly the same when switching from TFP to LAP. They contribute 14% (or 8 percentage points) to growth in aggregate LAP (1997–2007). Furthermore, inter-industry effects appear to be most pronounced during periods of economic downturns (as highlighted by the grey areas in Figure 3), similar to the developments found for TFP.

Figure 3: LAP growth (left) and decomposed intra-industry effect (right)

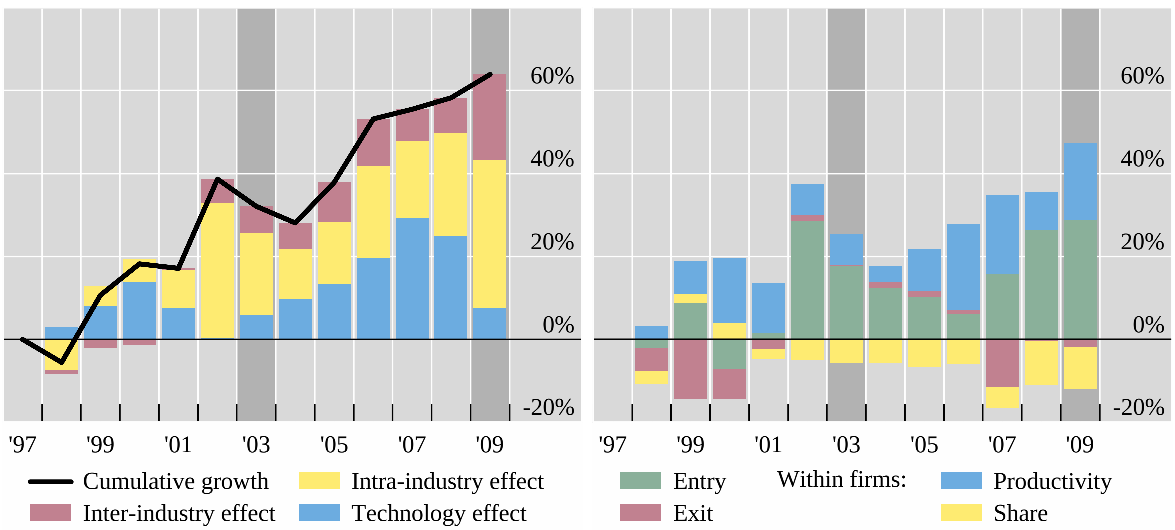

Decomposing the intra-industry effect into its four components (Figure 3, right-hand panel) unveils some noticeable differences in the results obtained for LAP and TFP. Most importantly, we observe a significant rise in the firm-entry effect when switching from TFP to LAP. One interpretation is that LAP underestimates the productivity advantage of new firms.[16] Small firms entering the market may appear to have below-average LAP, contributing to a positive firm-entry effect via the impact on the allocative efficiency. Assessing firm-level productivity based on TFP, however, takes account of firms' use of all inputs and can help correct for the underestimation of productivity. Small, but highly productive, entrants would thus tend to reduce the entry effect, notwithstanding that in both cases the impact on the average productivity of entering firms will mitigate the effect. That said, the impact of identification challenges, as discussed in Section 4.2, warrant some caution in interpreting the results.

[16] As discussed in footnote 15, revenue-based TFP measures may underestimate the productivity advantage as well. In this sense, the downward bias in LAP may thus be even more pronounced than for TFP.

To complete the comparison of TFP and LAP results, we present the contribution of each of the three main effects to the yearly variation in aggregate LAP in columns three and four of Table 3, using sales- and employment-based weights, respectively. Intra-industry effects remain the most important driver of yearly changes, whereas reallocations across industries seem to have no statistically significant explanatory power. While the role for changes in the productivity of continuing firms is surprisingly small for LAP when using sales weights (third column), the results of the variance decomposition for LAP based on employment shares (fourth column) are relatively closely aligned with those from our TFP estimates.

5. Conclusion

In this paper, we develop a framework to assess the contributions to aggregate productivity growth from inter- and intra-industry reallocations of resources as well as from technological progress. The framework quantifies the importance of these three channels, each of which has been emphasised in different strands of the international trade literature. Furthermore, we decompose the intra-industry effect into four distinct elements that feature in the new trade theory, ie the within-firm productivity, within-firm share, firm-entry and firm-exit effects. The firm-entry and firm-exit effects are each composed of two elements which makes it clear that entering and exiting firms affect aggregate productivity not only through their average productivity compared with continuing firms but also through an allocation effect that includes the relative size and productivity of each of these firms.

An empirical example based on Swiss manufacturing firms' TFP from 1997 to 2009 illustrates how the framework can be implemented and what conclusions can be drawn from the measures it generates. Our empirical application suggests that intra-industry reallocations among heterogeneous firms, as emphasised by Melitz and Trefler (2012), are an important source of aggregate productivity growth. In our study of Swiss manufacturing, the intra-industry effect accounts for more than half of the growth in the sector's TFP (1997–2007). Among the different effects contributing to intra-industry reallocations, productivity growth within large firms as well as firm entry represent key factors. Our analysis also implies that *inter-industry specialisation*, as underscored by (neo-)classical trade theory, remains important. Reallocations of resources towards relatively more productive industries account for about 11% of the increase in aggregate TFP. All these effects are found for a period characterised by the negotiation and introduction of the Swiss-EU Bilateral Agreements I and II and the increasing economic integration of the Swiss economy.

A more detailed study of the dynamics during recessions would provide an interesting field for further research. For one, we find the inter-industry effect to have been particularly strong during the two recessions observed in our sample, whereas changes in the technology effect differ markedly. This raises the question whether specific firm characteristics, such as financial leverage and international integration, can help explain differences in industries' adjustments to these adverse shocks. Extending our framework to account for the effect of firms switching industries or operating in multiple industries, none of which is the case in our sample, could also open up the perspective on additional adjustment channels.

Our analysis generally implies that the different strands of trade literature – the traditional and new trade theories – should be regarded as complements in explaining changes in aggregate productivity specifically, or the gains from trade in general. In addition, the decomposition of the intra-industry effect provides a helpful basis for disentangling the growth effects discussed in the new trade theory. More research is needed, however, to discriminate between these effects and their interplay.

References

Ackerberg, D. A., Caves, K., Frazer, G., 2006. Structural Identification of Production Functions, mimeo.

Aghion, P., Howitt, P., March 1992. A model of growth through creative destruction. Econometrica 60 (2), 323–351.

Atalay, E., June 2014. Materials prices and productivity. Journal of the European Economic Association 12 (3), 575–611.

Baily, M. N., Hulten, C., Campbell, D., Bresnahan, T., Caves, R. E., 1992. Productivity Dynamics in Manufacturing Plants. Brookings Papers on Economic Activity. Microeconomics 1992, 187–267.

Bartelsman, E. J., Doms, M., 2000. Understanding Productivity: Lessons from Longitudinal Microdata. Journal of Economic Literature 38 (3), 569–594.

Bernard, A. B., Eaton, J., Jensen, J. B., Kortum, S. S., 2003. Plants and Productivity in International Trade. American Economics Review 93 (4), 1268–1290.

Bernard, A. B., Jensen, J. B., Redding, S. J., Schott, P. K., 2007. Firms in International Trade. Journal of Economic Perspectives 21 (3), 105–130.

Bombardini, M., Kurz, C., Morrow, P., 2012. Firm Heterogeneity and Ricardian Trade: The Impact of Domestic Competition on Export Performance. Canadian Journal of Economics 45 (2), 585–612.

Brülhart, M., 2000. Dynamics of Intraindustry Trade and Labor-Market Adjustment. Review of International Economics 8 (3), 420–435.

Bustos, P., 2011. Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms. American Economic Review 101 (1), 304–340.

Coe, D. T., Helpman, E., Hoffmaister, A. W., 2009. International R&D spillovers and institutions. European Economic Review 53, 723–741.

Costinot, A., Donaldson, D., May 2012. Ricardo's Theory of Comparative Advantage: Old Idea, New Evidence. American Economic Review Papers and Proceedings 102 (3), 453–458.

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., Konings, J., 2006. Job reallocation and productivity growth in a post-socialist economy: Evidence from Slovenian manufacturing. European Journal of Political Economy 22, 388–408.

Debaere, P., 2003. Relative Factor Abundance and Trade. Journal of Political Economy 111 (3), 589–610.

Demidova, S., 2008. Productivity Improvements and Falling Trade Costs: Boon or Bane? International Economic Review 49, 1437–1462.

Eaton, J., Kortum, S., September 2002. Technology, Geography, and Trade. Econometrica 70 (5), 1741–1779.

Farinas, J. C., Ruano, S., 2005. Firm Productivity, Heterogeneity, Sunk Costs and Market Selection. International Journal of Industrial Organization 23, 505–534.

Finicelli, A., Pagano, P., Sbracia, M., 2013. Ricardian Selection. Journal of International Economics 89, 96–109.

Foster, L., Haltiwanger, J., Krizan, C. J., 2001. Aggregate Productivity Growth. Lessons from Microeconomic Evidence. In: Hulten, C. R., Dean, E. R., Harper, M. J. (Eds.), New Developments in Productivity Analysis. University of Chicago Press, Ch. 8, pp. 303–372.

Foster, L., Haltiwanger, J., Krizan, C. J., 2006. Market Selection, Reallocation, and Restructuring in the U.S. Retail Trade Sector in the 1990s. The Review of Economics and Statistics 88 (4), 748–758.

Foster, L., Haltiwanger, J., Syverson, C., 2008. Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? American Economic Review 98 (1), 394–425.

Greenaway, D., Kneller, R., 2007. Firm Heterogeneity, Exporting and Foreign Direct Investment. Economic Journal 117 (517), F134–F161.

Greenaway, D., Upward, R., Wright, P., 2000. Sectoral transformation and labour-market flows. Oxford Review of Economic Policy 16 (3), 57–75.

Görg, H., Hanley, A., Strobl, E., 2008. Productivity Effects of International Outsourcing: Evidence from Plant-level Data. Canadian Journal of Economics 41 (2), 670–688.

Grilliches, Z., Regev, H., 1995. Firm Productivity in Israeli Industry: 1979-1988. Journal of Econometrics 65, 175–203.

Grossman, G. M., Helpman, E., 1991. Innovation and Growth in the Global Economy. MIT Press, Cambridge.

Keller, W., 2004. International Technology Diffusion. Journal of Economic Literature 42 (3), 752–782.

Klette, T. J., Griliches, Z., 1996. The Inconsistency of Common Scale Estimators when Output Prices are Unobserved and Endogenous. Journal of Applied Econometrics 11 (4), 343–361.

Krugman, P. R., 1979. Increasing Returns, Monopolistic Competition, and International trade. Journal of International Economics 9 (4), 469–479.

Krugman, P. R., 1980. Scale Economies, Product Differentiation, and the Pattern of Trade. American Economic Review 70 (5), 950–959.

Levinsohn, J., Petrin, A., 2003. Estimating Production Functions Using Inputs to Control for Unobservables. Review of Economic Studies 70 (2), 317–341.

Lööf, H., Andersson, M., 2010. Imports, Productivity and Origin Markets: The Role of Knowledge-intensive Economies. The World Economy 33 (3), 458–481.

Lileeva, A., Trefler, D., 2010. Improved Access to Foreign Markets Raises Plant-Level Productivity...For Some Plants. Quarterly Journal of Economics 125 (3), 1051–1099.

Mairesse, J., Jaumandreu, J., 2005. Panel-data Estimates of the Production Function and the Revenue Function: What Difference Does It Make? Scandinavian Journal of Economics 107 (4), 651–672,.

Melitz, M. J., 2003. The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. Econometrica 71 (6), 1695–1725.

Melitz, M. J., Polanec, S., June 2012. Dynamic Olley-Pakes Productivity Decomposition with Entry and Exit, mimeo.

Melitz, M. J., Trefler, D., 2012. Gains from Trade when Firms Matter. Journal of Economic Perspectives 26 (2), 91–118.

Olley, G. S., Pakes, A., 1996. The Dynamics of Productivity in the Telecommunications Equipment Industry. Econometrica 64 (6), 1263–1297.

Ornaghi, C., 2006. Assessing the Effects of Measurement Errors on the Estimation of Production Functions. Journal of Applied Econometrics 21 (6), 879–891.

Petrin, A., White, T. K., Reiter, J. P., 2011. The impact of plant-level resource reallocations and technical progress on U.S. macroeconomic growth. Review of Economic Dynamics 14, 3–26.

Romer, P. M., 1990. Endogenous Technological Change. Journal of Political Economy 98, S71–S102.

Schröder, P. J. H., Sørensen, A., 2012. Firm exit, technological progress and trade. European Economic Review 56, 579–591.

Shin, K., 1997. Inter- and Intrasectoral Shocks: Effects on the Unemployment Rate. Journal of Labor Economics 15, 376–401.

Van Beveren, I., 2012. Total Factor Productivity Estimation: A Practical Review. Journal of Economic Surveys 26 (1), 98–128.

Van Biesebroeck, J., 2007. Robustness of Productivity Estimates. Journal of Industrial Economics 55, 529–569.

Van Biesebroeck, J., 2008. The Sensitivity of Productivity Estimates. Journal of Business & Economic Statistics 26, 311–328.

Wagner, J., 2007. Exports and Productivity: A Survey of the Evidence. The World Economy 30 (1), 60–82.

Wagner, J., 2010. Entry, Exit and Productivity: Empirical Results for German Manufacturing Industries. German Economic Review 11 (1), 78–85.

Appendix

A. Derivation of the decomposition in equation (4)

We start from the definition of aggregate productivity, Φ_t, in equation (2), denoting the difference in productivity from period $t = 0$ to t as $\Phi_t - \Phi_0$. Adding and subtracting both, $\sum_j \sum_i S_{j0}\Delta s_{ijt}\Delta\varphi_{ijt}$ and $\sum_j S_{j0}\bar{\varphi}_{jt}$, to this difference and collecting terms yields:

$$\begin{aligned}\Phi_t - \Phi_0 \;=\; & \sum_j \sum_i \left(S_{jt} - S_{j0}\right)\Delta s_{ijt}\Delta\varphi_{ijt} + \sum_j \left(S_{jt} - S_{j0}\right)\bar{\varphi}_{jt} \\ & + \sum_j S_{j0}\left(\sum_i \Delta s_{ijt}\Delta\varphi_{ijt} - \sum_i \Delta s_{ij0}\Delta\varphi_{ij0}\right) + \sum_j S_{j0}\left(\bar{\varphi}_{jt} - \bar{\varphi}_{j0}\right). \qquad (5)\end{aligned}$$

To derive the inter-industry effect, we summarise the first two terms on the right-hand side of equation (5) as:

$$\sum_j \left(S_{jt} - S_{j0}\right)\left[\left(\sum_i \Delta s_{ijt}\Delta\varphi_{ijt}\right) + \bar{\varphi}_{jt}\right] \;=\; \sum_j \left(S_{jt} - S_{j0}\right)\Phi_{jt}. \qquad (6)$$

Note from equation (2) that the term in squared brackets is equal to $\Phi_{jt} = \sum_i (s_{ijt}/S_{jt})\varphi_{ijt}$ which is the industry-equivalent of aggregate productivity, since $s_{ijt} \,/\, S_{jt}$ corresponds to each firm's sales share within its industry. To facilitate the interpretation, we state the inter-industry effect as:

$$\sum_j \left(S_{jt} - S_{j0}\right)\Phi_{jt} = \sum_j \left(S_{jt} - S_{j0}\right)\Delta\,\Phi_{jt}, \qquad (7)$$

where $\Delta\,\Phi_{jt}$ corresponds to the deviation of Φ_{jt} from the unweighted mean of industry productivity levels, $\bar{\Phi}_t = J^{-1}\sum_{j\in J}\Phi_{jt}$. Equation (7) holds because $\sum_j \left(S_{jt} - S_{j0}\right)\bar{\Phi}_t = \bar{\Phi}_t \sum_j \left(S_{jt} - S_{j0}\right) = 0$, since the deviation of industry shares in t from their levels in $t-1$ must add up to zero.

To derive the four effects that collectively represent the intra-industry effect as well as the technology effect in equation (4), we exclude the inter-industry effect (the first two terms on the right-hand side of equation (5)) and write the other two remaining terms, now distinguishing between continuing (C), entering (E) and exiting firms (X), as:

$$\begin{aligned} & \sum_j S_{j0}\left(\sum_i \Delta s_{ijt}\Delta\varphi_{ijt} - \sum_i \Delta s_{ij0}\Delta\varphi_{ij0}\right) + \sum_j S_{j0}\left(\bar{\varphi}_{jt} - \bar{\varphi}_{j0}\right) \\ = \; & \underbrace{\sum_j \sum_{i\in C} S_{j0}\left(\Delta s_{ijt}\Delta\varphi_{ijt} - \Delta s_{ij0}\Delta\varphi_{ij0}\right)}_{\text{within-firm effects}} \\ & + \sum_j \sum_{i\in E} S_{j0}\,\Delta s_{ijt}\Delta\varphi_{ijt} - \sum_j \sum_{i\in X} S_{j0}\,\Delta s_{ij0}\Delta\varphi_{ij0} + \sum_j S_{j0}\left(\bar{\varphi}_{jt} - \bar{\varphi}_{j0}\right). \qquad (8)\end{aligned}$$

The first term on the right-hand side of (8) summarises the within-firm productivity and share effects, which are both based on developments of continuing firms only. In order to separate these two effects, we rewrite this first term to get:

$$\begin{aligned}\underbrace{\sum_j \sum_{i \in C} S_{j0} \left(\Delta s_{ijt} \Delta\varphi_{ijt} - \Delta s_{ij0} \Delta\varphi_{ij0}\right)}_{\text{within-firm effects}} &= \underbrace{\sum_j \sum_{i \in C} S_{j0} \, \Delta \, s_{ijt} \, (\Delta \, \varphi_{ijt} - \Delta \, \varphi_{ij0})}_{\text{within-firm productivity effect}} \\ &+ \underbrace{\sum_j \sum_{i \in C} S_{j0} \, \Delta \, \varphi_{ij0} \, (\Delta \, s_{ijt} - \Delta \, s_{ij0})}_{\text{within-firm share Effect}} . \qquad (9)\end{aligned}$$

The remaining three terms on the right-hand side of equation (8) measure the contribution of entering and exiting firms to the allocative efficiency of the manufacturing sector, respectively, and a general technology effect that includes all types of firm.

To develop the complete firm-entry and firm-exit effect and to restrict the technology effect to continuing firms, we decompose the term $\sum_j S_{j0} \left(\bar{\varphi}_{jt} - \bar{\varphi}_{j0}\right)$ in (8) by distinguishing between continuing and exiting firms in period $t = 0$ and continuing and entering ones in period t:

$$\begin{aligned}\sum_j S_{j0} \left(\bar{\varphi}_{jt} - \bar{\varphi}_{j0}\right) &= \sum_j S_{j0} s^E_{jt} \, \bar{\varphi}^E_{jt} + \sum_j S_{j0} \left(1 - s^E_{jt}\right) \bar{\varphi}^C_{jt} \\ &- \sum_j S_{j0} s^X_{j0} \, \bar{\varphi}^X_{j0} - \sum_j S_{j0} \left(1 - s^X_{j0}\right) \bar{\varphi}^C_{j0}, \qquad (10)\end{aligned}$$

where we make use of the definition of productivity averages for each firm category, $\bar{\varphi}^{\gamma}_{jt} = (\sum_{i \in \gamma} \varphi_{ijt})/N^{\gamma}_{jt}$ for $\gamma = \{C, E, X\}$. In addition, we define the share of the number of entering firms in the total number of active firms of the same period as s^E_{jt}. By analogy, we define the relative number of firms that exit in t but are still active in $t = 0$ to the total number of firms in $t = 0$ as s^X_{j0}. Rearranging terms then yields:

$$\begin{aligned}\sum_j S_{j0} \left(\bar{\varphi}_{jt} - \bar{\varphi}_{j0}\right) &= \underbrace{\sum_j S_{j0} \left(\bar{\varphi}^C_{jt} - \bar{\varphi}^C_{j0}\right)}_{\text{technology effect}} \\ &+ \sum_j S_{j0} s^E_{jt} \left(\bar{\varphi}^E_{jt} - \bar{\varphi}^C_{jt}\right) - \sum_j S_{j0} s^X_{j0} \left(\bar{\varphi}^X_{j0} - \bar{\varphi}^C_{j0}\right) . \qquad (11)\end{aligned}$$

The last two terms on the right-hand side of equation (11) represent the level impact on aggregate productivity from firm entry and exit, respectively. Substituting the results in equation (11) and (9) into equation (8) then leads to the technology effect and the four intra-industry effects shown in equation (4).

B. TFP estimation method

We assume a Cobb-Douglas production function, $Q_{ijt} = \varphi_{ijt} \, K^{\alpha_k}_{ijt} \, L^{\alpha_l}_{ijt} \, M^{\alpha_m}_{ijt}$, with α_z for $z = \{k, l, m\}$ representing the elasticity of substitution of each input factor. j identifies the firm's affiliation to an industry and t the time period. By taking natural logs (denoted as lower-case letters) the basic estimation equation for TFP becomes $q_{ijt} = \alpha_0 + \alpha_k k_{ijt} + \alpha_l l_{ijt} + \alpha_m m_{ijt} + \omega_{ijt} + \varepsilon_{ijt}$, where we assume the functional form for the firm's TFP to be $\varphi_{ijt} = \exp(\alpha_0 + \omega_{ijt} + \varepsilon_{ijt})$. The term α_0 is a time-invariant measure of average productivity across firms, whereas ω_{ijt} represents the firm-specific productivity that is observable to the firm and (potentially) predictable by the researcher. The last term, ε_{ijt}, encompasses a random shock to productivity at the firm level and any measurement errors. According to these widely used assumptions, ω_{ijt} is the only source of systematic firm heterogeneity.

In the following, we address three potential sources that may bias our TFP estimates: (i) omitting firm-level prices; (ii) the endogeneity of the firm's input choices; and (iii) sample selection.

Approximating physical output, q_{ijt}, by revenues deflated at the industry-level, $\tilde{r}_{ijt} = q_{ijt} + p_{ijt} - p_{jt}$ (where p_{ijt} and p_{jt} denote the logs of the firm's output price and the industry's price level, respectively), will bias the estimated coefficients in the production function if the variation in firms' output prices is correlated with the firms' input choices. Following Klette and Griliches (1996), we assume each firm i, operating in industry j, faces a CES demand structure for the variety it produces (in logs), $q_{ijt} = q_{jt} - \sigma_j\,(p_{ijt} - p_{jt}) + \upsilon_{ijt}$. The industry's output is given by q_{jt} whereas υ_{ijt} captures any independent and identically distributed (i.i.d.) demand shocks or measurement errors. In optimising, each firm sets an individual price, p_{ijt}, equal to its marginal cost of production times the industry-specific markup, $\sigma_j/(\sigma_j - 1)$, with $\sigma_j > 1$ representing the elasticity of substitution between product varieties of the same industry. Since we observe a firm's revenue rather than the quantity produced, we solve the demand function for the firm's price and replace physical output by the definition in the production function above. We collect terms to obtain:

$$\tilde{r}_{ijt} = \beta_0 + \beta_k\, k_{ijt} + \beta_l l_{ijt} + \beta_m\, m_{ijt} + \beta_j\, q_{jt} + \omega^*_{ijt} + \varepsilon^*_{ijt}, \tag{12}$$

with the coefficients of interest, $\beta_z = \alpha_z\,(\sigma_j - 1)/\sigma_j$ for $z = \{0, k, l, m\}$, and $\beta_j = 1/\sigma_j$. The latter allows the elasticity of substitution to be identified. Furthermore, $\omega^*_{ijt} = \omega_{ijt}(\sigma_j - 1)/\,\sigma_j$, and i.i.d. shocks are summarised in $\varepsilon^*_{ijt} = \varepsilon_{ijt}\,(\sigma_j - 1)/\,\sigma_j + \upsilon_{ijt}/\sigma_j$. The calculation of TFP estimates at the firm level thus requires the estimated markup in each industry to be taken into account:

$$\hat{\omega}_{ijt} + \hat{\alpha}_0 = \left(\frac{\hat{\sigma}_j}{\hat{\sigma}_j - 1}\right)\Big(\tilde{r}_{ijt} - \hat{\beta}_k\, k_{ijt} - \hat{\beta}_l\, l_{ijt} - \hat{\beta}_m\, m_{ijt} - \hat{\beta}_j\, q_{jt}\Big). \tag{13}$$

We retrieve industry-specific elasticities by replacing $\beta_j\, q_{jt}$ with $\sum_{j=1}^{J} \beta_j q_{jt} I_{ij}$, where I_{ij} is a dummy variable of unit value (otherwise zero) if firm i is active in industry j, to benefit from pooling efficiencies, rather than running the regression for each industry separately. While this setup allows for a firm to face up to J different demand functions, a single firm's activity is always assigned to exactly one industry in our data set.

To ensure unbiased estimations for the firms' TFP, we follow the procedure outlined by Ackerberg et al. (2006). They impose an explicit timing assumption. Each firm decides whether to invest or not at $t-1$, thereby fixing the level of available capital at t, k_{ijt}. In the intermediate period $t-b$ with $0 < b < 1$, the firm chooses the amount of labour it will employ at t, l_{ijt}. Finally, at time t, the firm selects the intermediates it uses, m_{ijt}. Productivity evolves according to a first-order Markov process between these time periods. We thus have $\omega_{ijt-b} = E\left(\omega_{ijt-b}|\omega_{ijt-1}\right) + \xi_{ijt-1}$ and $\omega_{ijt} = E\left(\omega_{ijt}|\omega_{ijt-b}\right) + \xi_{ijt-b}$ with ξ_{ijt} representing the innovation to the firm's productivity at time t . This implies that firms base their expectations about future productivity only on given productivity levels. The firm's investment decision in t, i_{ijt}, will therefore rely on the firm's given capital stock, k_{ijt}, contemporaneous labour employed, l_{ijt}, as well as the most recently observed productivity level ω_{ijt}, so that we can refer to the investment function as $i_{ijt} = i_t(k_{ijt},\, l_{ijt},\, \omega_{ijt})$. Assuming that investment is strictly monotonic in productivity, we invert the investment function to obtain $\omega_{ijt} = h_t(k_{ijt},\, l_{ijt},\, i_{ijt})$.

The remainder of the estimation approach relies on Olley and Pakes (1996). This involves approximating $h_t(\cdot)$ with a second-order polynomial in investment, the proxy variable, as well as capital and labour, the state variables. Embedding $h_t(\cdot)$ in an OLS regression of (12) then yields consistent estimates of the coefficients of the industry's elasticity of substitution as well as of intermediates which the firm chooses after observing the productivity change, $\tilde{r}_{ijt} = \beta_m m_{ijt} + \beta_j q_{jt} + \phi_t(k_{ijt}, l_{ijt}, i_{ijt}) + \varepsilon^*_{ijt}$, where $\phi_t(k_{ijt}, l_{ijt}, i_{ijt}) = \beta_0 + \beta_k k_{ijt} + \beta_l l_{ijt} + [(\sigma_j - 1)/\sigma_j] h_t(k_{ijt}, l_{ijt}, i_{ijt})$. Due to the non-parametric treatment of $h_t(\cdot)$ both β_k and β_l, remain unidentified in the OLS regression. These coefficients are recovered by exploiting the dynamics of the productivity change as well as information from the firm's exit decision in a second step.

This second step consists of non-parametrically regressing ω_{ijt} on ω_{ijt-1} to obtain coefficient estimates of capital and labour. Given the assumed Markov process, the expected value of the productivity shock, conditional on the firm being active in t can be expressed as $E\left(\omega_{ijt}\,|\,\omega_{ijt-1}, \chi_{ijt} = 1\right) = g\left(P_{ijt},\, \omega_{ijt-1}\right)$ where χ_{ijt} represents an indicator variable that takes zero value if the firm exits in period t and unit value otherwise. We use the latter to fit a probit model on a polynomial of the first lags of the state and proxy variables to obtain estimates of the probability of survival, $\hat{P}_{ijt}$. The unknown function $g(\cdot)$ is approximated by a second-order polynomial in $\hat{P}_{ijt}$ and $(\hat{\phi}_{t-1} - \beta_k\, k_{ijt-1} - \beta_l\, l_{ijt-1})$, with $\hat{\phi}_{t-1}$ representing the estimate of ϕ_{t-1} using the results from the initial OLS regression. Using the estimated coefficients on intermediates and industry production from the first stage OLS regression, a consistent estimate of the capital and labour coefficient follows from a non-linear regression of:

$$\begin{aligned} \tilde{r}_{ijt} - \hat{\beta}_m\, m_{ijt} - \hat{\beta}_j\, q_{jt} \quad &= \quad \beta_0 + \beta_k\, k_{ijt} + \beta_l\, l_{ijt} \\ &\quad + g\left(\hat{P}_{ijt},\, \hat{\phi}_{t-1} - \beta_k k_{ijt-1} - \beta_l l_{ijt-1}\right) + \xi_{ijt} + \varepsilon^*_{ijt}. \end{aligned} \tag{14}$$

Once all coefficient estimates are retrieved, TFP estimates follow from (13). Standard errors are obtained from bootstrapping.