Slowing Down or Racing Ahead: Heterogeneity in Firm-performance in the Netherlands

Michael Polder

(Statistics Netherlands)

Paper prepared for the 35th IARIW General Conference

Copenhagen, Denmark, August 20-25, 2018

Session 4A: Is Productivity Growth Slowing Down? I

TUESDAY, AUGUST 21, AFTERNOON

Slowing down or racing ahead? Heterogeneity in firm-performance in the Netherlands

Michael Polder

~ preliminary and not for quotation ~

Abstract

Using firm-level data for the Netherlands, we explore whether the productivity slowdown is an economy wide phenomenon, or whether there is firm heterogeneity in this respect. We calculate aggregate productivity patterns from firm-level data, and look at measures of dispersion. In addition, we investigate the persistence of productivity performance, and document differences in firm characteristics across the productivity distribution. Finally, we investigate the correlation of worker income with firm-level productivity. The paper presents a first, preliminary set of results, and at this stage is not for quotation.

1. Introduction

The joke about the statistician that drowned in the water of an on average shallow river has been told too many times. As with many clichés, however, there is a core of truth on which it is based. Apart from breakdowns by classes, official statistics rarely give insight into dispersion, or any other distributional features, especially macro-economic figures provided by the National Accounts. This is true also for productivity statistics, which usually do not go beyond a breakdown by industry. In this sense, the productivity slowdown that is evident in most developed countries, from at least the beginning of the new century, can be qualified as an aggregate phenomenon. Firm-level research on productivity, going back at least three decades by now, has convincingly documented the dispersion in productivity performance across firms, even in narrowly defined industries (see Bartelsman and Doms, 2000, for an early survey of the literature; and Syverson, 2011 for an update). Only very recently have statistical offices begun to use these insights to complement their productivity statistics with information on the underlying firm heterogeneity, see Cunningham et al. (2018) for the US, and Ardanaz-Badia et al. (2017) for the UK.

A slowdown in the growth of productivity has been documented for most developed countries (OECD, 2017a). While the crisis has had a strong impact on economic performance, the slowdown has started well before that, and therefore it cannot be the sole reason for the meagre productivity figures. From a macro-economic point of view, explanations offered include the lack of innovation and technological change (e.g. Gordon, 2016) and secular stagnation (Summers, 2016), while the role of (mis)measurement in this respect seems to be small (Byrne et al. 2016).

As mentioned above, from the empirical literature on productivity making use of firm-level data, it is well known that there is substantial heterogeneity across firms with respect to the level and growth of productivity. OECD (2015, 2016) find that there is an increasing gap between firms at the production frontier and those that are behind, translating into an overall decrease in productivity performance. In fact, according to this work, productivity at the global frontier has been growing steadily, while growth for the lagging firms has been only marginal.[1] This finding could suggest a lack of diffusion of knowledge and technology from frontier firms to the rest of the economy. Moreover, evidence from the US suggests that the rate of creative destruction through business dynamics (i.e. entry, exit and reallocation) seems to have gone down (see Decker et al. 2018). This suggests that another

[1] Recent research for the Netherlands (Heuvelen et al. 2018) was unable to confirm this finding, although the time-period used in that study is shorter and more recent, and only the national frontier has been investigated.

channel affecting the slowdown might be the less efficient allocation of factors of production across firms.

Such differences in firm productivity are interesting not only to explain macro-economic performance, it can have implications for the wages that a firm can pay its employees as well (Mortensen, 2003). In this respect, OECD (2017b) finds evidence that firms with a higher productivity, pay higher wages to their employees. Productivity differences between firms can therefore be a concern for growing (income) inequality.

Our paper will look at the firm-level distribution of productivity (both labour and multifactor productivity) in the Netherlands. Our first goal is to document the heterogeneity in productivity performance surrounding the macro and industry-level figures from around the beginning of the century. Are there (growing) differences between frontier and non-frontier firms? The second goal is to compare the characteristics of firms on the frontier and those that are not. Are frontier firms more innovative (in a technological or non-technological sense), are they more ICT and research intensive, do they have a higher-quality labor force? Are there differences in their export behavior or type of ownership? The final goal is to document the relation between differences in productivity and wages, and assess whether heterogeneity in firm performance is a main driver of wage inequality.

The plan of the paper reads as follows. In next section we will outline a simple conceptual framework to explain the mechanisms through which changes in the firm-level distribution affect aggregate productivity growth. In section 3 we will briefly discuss related literature, which at this stage is a placeholder only. Section 4 describes the data, and section 5 the concept of productivity employed, and several matters of definitional and methodological order. Section 6 then documents the trends in productivity growth for the Netherlands, comparing different industries and also comparing to other countries. The results from the firm-level data are presented in section 7. First, we document the firm-level heterogeneity, the role of frontier of and non-frontier firms, persistence, and candidate determinants for explaining differences between firms. Finally, we present tentative evidence for the relation between firm-level productivity and wages. Section 8 concludes. While providing a first sketch of the research, we are aware that all this needs considerable more work.

2. Simple conceptual model

Firms make all kinds of decisions to optimize their performance, such as investment in physical capital, hiring and firing of workers, whether to become active in the international market through exporting, to source inputs domestically or abroad, to do R&D or innovate, et cetera. None of these decisions will necessarily lead to success in a deterministic sense. Rather, there is a lot of uncertainty surrounding these decisions, both in the information set available to the firm at the time of decision making, as well as about how the future will evolve, for example in terms of demand conditions, actions of competitors, and changes in the institutional environment.

A simple way to formalize the above is to view a firm's productivity as stochastic, i.e. a draw from a probability distribution, say $G(A)$. Productivity may be influenced by firm decisions and other external determinant X, which can be formulated as a conditional distribution, $G(A|X)$. Syverson (2011) notes that in equilibrium only firms that receive a draw that is favorable enough to overcome a fixed cost of entry will enter the market, so that the observed distribution is a truncated version of $G(A)$, subject to lower threshold $\underline{A}$.

Along this line of thought, a slowdown in productivity growth can arise for several reasons. For instance, productivity growth can result from a gradual overall shift of the productivity distribution to the right. When overall economic conditions turn less favorable, firms will receive lower productivity on average, and all else equal, this will result in lower overall productivity growth. Such a scenario can for example arise in case of the headwinds described by Gordon (2016), such as a declining demand due to ageing and globalization, stalling of increases in the educational level or worker skills, and lesser advances in technological progress.

In addition, lower aggregate productivity growth may result from a reallocation of inputs to less productive firms. That is, as aggregate growth is a weighted sum of productivities at the firm-level, where larger firms receive more weight, lower productivity can arise just by moving resources from the higher productive firms to the lower productive firms. Aggregate productivity as the weighted sum of firm-level can be written as

(1) $A_t = \sum_i w_{it} A_{it}$

Where w_{it} is the size weight of the firm which should be in line with the denominator use in the productivity definition. A decomposition popularized by Olley and Pakes (1996) shows that aggregate productivity can be written as an unweighted mean and the covariance of productivity with firm size

$$A_t = \overline{A}_t + \sum_i (w_{it} - \overline{w}_t)(A_{it} - \overline{A}_t)$$

Where $\overline{A}$ and $\overline{w}$ are simple averages across firms. The gap between the two is frequently used as measure of allocative efficiency, see e.g. Bartelsman et al. (2009). The idea is that, in this static framework, productivity could be enhanced by allowing more productive firms to grow. Recently, using a different conceptual model, issues of misallocation have received quite some attention through the work of Hsieh and Klenow (2009) and others.

The previous two scenarios do not involve a change in the shape of the productivity distribution. Changes in higher moments however can also affect the resulting overall growth. By way of illustration, consider a situation with high productive firms and low productive firms. Productivity growth may then occur if the low productive firms catch up to the level of the high productive firms. By consequence the productivity distribution will narrow. Conversely, the high productive firms may increase their productivity further, for example through innovation. In such a situation, the productivity gap will widen. In the catching-up scenario, a slowdown could be the consequence of the productivity gains of the laggards not being able to compensate the lower growth at the frontier. Vice versa, if the frontier firms improve their productivity, these gains may not be high enough to offset the lower off-frontier growth, and this will result in a slowing down of the aggregate productivity growth. These shifts could be related to differences in engagement in underlying productivity enhancing actions by individual firms, causing a change in the shape of the conditional distribution $g(A|X)$.

Another way through which productivity changes occur is through changes in the composition of the distribution. Clearly, the population of firms in the economy is not constant. In an efficient economy, less productive incumbent firms in the lower end of the distribution exit the market, to be replaced by more productive entrants, resulting in growth through Schumpeterian creative destruction. It follows that a decline in business dynamism can therefore have an adverse effect on productivity. Evidence of such a decline has been documented for the U.S. by Decker et al (2018), amongst others. In such a scenario, the contribution of successful startups bringing new ideas to the market becomes weaker, while inefficient incumbent firms are allowed to stay longer in the market, holding back overall growth. The contribution of entry and exit can be quantified through a decomposition of the growth in total productivity. Differencing equation (1) and separating out the contributions of continuing firms (C), entry (N) and exit (X), we get

(2) $$\Delta A_t = \sum_i w_{it} A_{it} - \sum_i w_{it-1} A_{it-1}$$

$$= \left(\sum_{i \in N} w_{it} A_{it} + \sum_{i \in C} w_{it} A_{it}\right) - \left(\sum_{i \in X} w_{it-1} A_{it-1} + \sum_{i \in C} w_{it-1} A_{it-1}\right)$$

$$= A_t(N) - A_{t-1}(X) + \Delta A_t(C)$$

Further, the contribution of continuing firms can be rewritten into the contribution of (within-firm) productivity changes, and the contribution of changes in (between-firm) resource allocation

$$\Delta A_t(C) = \sum_{i \in C} \Delta A_{it} w_{it-1} + \sum_{i} A_{it-1} \Delta w_{it}$$
$$= \Delta A_t^w(C) + \Delta A_t^b(C)$$

Clearly, productivity growth benefits from positive within and between firm changes, $\Delta A_t^w(C) > 0$ and $\Delta A_t^b(C) > 0$, as well as a positive balance between the contribution of entry and exit, $A_t(N) - A_t(X) > 0$.

Finally, the scenarios described above are not mutually exclusive. It may well be possible that less favorable economic circumstances withhold average growth, while at the same time institutional barriers prohibit resources to flow to the most efficient firms, and in addition there is an increase in the productivity gap between frontier and non-frontier firms, the latter being replaced by more efficient startups only to a low degree. This paper will try to present descriptive evidence highlighting the role of these firm-level forces behind aggregate productivity growth in the Netherlands. In this version, the results are preliminary, and at some points there is a gap between the conceptual model and the results presented.

3. Literature

(This is a place holder. A more extensive literature review is pending. For now, we mention that our work is closely related to ongoing work at the OECD (2015, 2017), and can be seen as a complementary exercise, looking more closely at the Netherlands. Other evidence on the micro-level sources of productivity growth comes from Lopèz-Garcia (2018, forthcoming). CPB (2018) has recently also looked at productivity patterns in the firm-level data. Bartelsman (2000) and Syverson (2011) survey the international literature on determinants of productivity. Faggio et al. (2007) have documented the relation between firm-level productivity and wages.

4. Data

Our macro-economic figures for the Netherlands are based on the Dutch Growth accounts (link), which span the period from 1995 to 2016.[2] Following the official figures published by Statistics Netherlands we will follow a model with an exogenous rate of return, thereby allowing for non-zero profits (for more detail, see De Haan et al 2014). For a better comparison to the micro-aggregated figures we exclude the financial sector (NACE Rev. 2 Chapter K), for which we do not

[2] Data for 2015 and 2016 are preliminary. Following the revision of the Dutch National Accounts this year, the Growth Accounts will also be revised in the coming period.

have firm-level data. For international comparisons we use the OECD productivity database (link), which is also based on the assumption of an exogenous rate of return. Note that the figures for the Netherlands are comparable, but not exactly equal, as a consequence of for instance differences in the derivation of capital services (e.g. service life of capital goods), and a different industry coverage (OECD covers the total economy, whereas the Statistics Netherlands data cover the so-called commercial sector, which excludes public administration, education, households, and real estate). Industry deflators for value added and capital inputs are sourced from National Accounts, and contain price information for 38 industries. For capital, we use the deflator of the total capital stock.

Firm-level information is sourced from the statistics on non-financial enterprises (NFO). This source contains information at the level of the enterprise group (@@@ link Eurostat), which is a consolidation of the enterprise level. For most observations, the enterprise and the enterprise group coincide. Typically, however larger firms may consist of a group of more than one enterprise. In general, the enterprise is a statistical unit that can be thought of as homogeneous in terms of economic activity and which can be described as a separate entity. The enterprise group can be thought of as the actor in the financial process, possibly comprising more than one economic activity (horizontal or vertical integration). For example, a retail firm may be active in the supermarket sector, but also own other more specialized shops, such as drugstores and liquor stores. A tech firm may be the owner of one or more dedicated R&D units. In the NFO data we observe the consolidated information for such holdings. An advantage of the NFO data is that it has a high coverage of smaller firms. In principle, all domestics firms liable to pay profit tax in the Netherlands are covered. This advantage comes at the cost of a more aggregated representation of the larger firms, through which a certain degree of firm heterogeneity is lost, as is clear from the example of the retail firm. On the other hand, the consolidation of the dedicated R&D unit into the overall firm entity is not necessarily bad, given that firms have some freedom in deciding where profits are recorded, thereby troubling the measurement of productivity at the more disaggregated level of the business unit (i.e. the enterprise rather than the enterprise group).

We use the Business Register (ABR) for employment information at the level of the enterprise. Given that the ABR contains the population of firms in the Netherlands, we also use it to construct sample weights when we want to construct aggregate measures. In addition, the ABR contains information on reasons for entry and exit of firms in and out of the population. This allows to distinguish actual births and death of firms, from merger and acquisitions, reorganization of units, or statistical reasons for renumbering firm units.

For making a link to firm characteristics, we will calculate productivity from the Production Statistics, containing information about the production structure of firms, including value added at factor cost, full-time equivalent employment, and depreciation cost as a proxy for capital inputs. All variables are according to Structural Business Statistics regulation of the European Commission. This source is in principle restricted to firms with more than 20 employees. However, it has the advantage, compared to the NFO data, that is can readily be linked to other surveys. We will use information from the Community Innovation Survey (CIS), the Investment Survey, and the Foreign Affiliates Trade Statistics (FATS).

5. Methodology and definitions

5.1 Productivity

Productivity is measured as output over input. We will use a value added based model, with capital and labor as the only inputs into production (KL-VA model).[3] We assume that output Y, capital inputs K and labor inputs L are homogeneous within firms. All variables are expressed in real terms (L will be operationalized as full-time equivalents jobs, fte). Labor productivity is defined as

$$LP = Y/L$$

TFP growth can be defined as the volume change (Q) of output over the volume change of total inputs (KL)

(1) $$TFPG = \frac{QY}{QKL}$$

Statistics Netherlands uses a Laspeyres type index to aggregate the inputs:

$$QKL = w_K QK + w_L QL$$

where weights w are defined as cost shares in the previous year.

In the firm-level analysis, besides productivity growth, we will also be interested in determining the productivity level. Typically, however, capital services are not directly observed in firm-level data. Using a proxy, of which it is assumed that it is proportional to actual capital services, it is convenient to use a Cobb-Douglas type aggregator. The advantage is that any proportional errors cancel out when growth figures are calculated, and the same holds when comparing levels across firms. Then denoting log transformed variables in lower case:

(2) $$tfp = y - w_K k - w_L l$$

In this case, under the assumption of constant returns to scale and no adjustment costs, cost shares equal the output elasticities. To stay as close as possible to the aggregate growth accounting figures, we will use the (lagged) industry-level cost

[3] It is also possible to base productivity on gross output and capital, labor and intermediate inputs as inputs into production (Y-KLEMS model). The KL-VA approach however seems more in line with the macro-economic evidence on the productivity slowdown, where GDP per capita or worker, and industry value added per unit of labor, are commonly used. Moreover, aggregation of gross output based measures of productivity is more difficult at the firm-level because of the issue of consolidation.

shares as weights in the firm-level productivity figures as well.[4] Note further that the TFP growth corresponding to the Cobb-Douglas levels (2) is approximately equal to the Laspeyres TFP growth for growth rates around 0, using that $\ln(QX) \approx QX - 1$. Finally, note that the factor weights w are time-dependent. To align with growth accounting definition, whenever calculating TFP growth, we use t-1 weights for both t-1 and t

$$(3) \qquad \begin{aligned} \Delta tfp &= tfp_1 - tfp_0 \\ &= \left(y_1 - w_{K,0}k_1 - w_{L,0}l_1\right) - \left(y_0 - w_{K,0}k_0 - w_{L,0}l_0\right) \\ &= \Delta y - w_K \Delta k - w_L \Delta l \end{aligned}$$

Since equation (2) uses lagged weights, when we use productivity levels, we are referring to tfp_1. In this definition, the level of TFP is determined up to an arbitrary time-invariant constant. For this reason, we normalize the firm's TFP level with the industry average.[5] This normalization also allows to compare firms across industries.

5.1.1 Frontier firms

Frontier firms are defined as those firms that are in the top of the productivity distribution. Defining the "top" involves a sense of arbitrariness. We shall consider the top decile as the best performing firms (either in terms of labor productivity or TFP, depending on the analysis). In addition, to avoid the frontier to be sensitive to incidental jumps in revenue or factor inputs, a degree of consistency is required. We therefore require firms to be active at the frontier for three consecutive years. With this definition, frontier firms are derived by industry, and by year. Therefore, in each industry 10% of the firms will be labelled as a frontier firm each year.

5.1.2 Weights

To be able to make a connection between the patterns observed at the micro-level, and those at the macro and industry-level, we have to make two corrections to the firm-level data. First, the firm-level data sources use for the productivity calculations are sample based. To correct for this, we use sampling weights based on populations counts from the Business Register. By industry and size-class, the number of firms in the population can be obtained, and compared to the number

---

[4] There have been significant advances in the econometric estimation of the production function, especially taking into account the endogeneity of factor input decisions when determining the output elasticities, see e.g. Olley and Pakes (1996); Levinsohn and Petrin (2003); Ackerberg et al. (2006); Wooldridge (2009). As noted by Melitz and Polanec (2012) this unlikely to significantly affect the productivity ranking of firms.

[5] Another way of thinking about this normalization is that the firm's level of productivity is benchmarked against that of another, fictive firm in the same industry, which has constant productivity equal to the average TFP level of all firms in that industry and across time. We do not use the industry average *by year* for the normalization, because that would affect our conclusions about the slowdown.

observed in the sample. The inverse sampling probability can be used as a weight to obtain a representative aggregate moment. Second, aggregate productivity figures can be seen as a size-weighted average of productivity at the firm-level. Thus, to mimic the aggregate pattern, firm-level productivity should be weighted by firm-size. In the case of labor productivity, one can use (lagged) employment. In the case of TFP, ideally one would use the total cost of production. Lacking the actual level of capital inputs (we use a proxy in the productivity calculations), this is not possible. Alternatively, output can be used. Given that value added can be negative, we will use turnover.

Three types of statistics can then be derived from the micro-data: 1. Unweighted moments which apply only to the pertinent sample; 2. Sample-weighted moments which are representative for the firms in the population (e.g. average productivity across firms in certain industry); 3. Sample-weighted and size-weighted moments, which should be representative for the aggregate (e.g. industry productivity growth). Note however, that even while the firm-level data are carefully weighted, they will never be exactly equal to the aggregate published totals from the National Accounts, as the latter figures result from integration of various sources, and a complex statistical process. Moreover, it is not possible with our micro-data to estimate productivity for self-employed, which are of course implicitly included in the National Accounts aggregates.

6. Macro and industry-level trends

In this section we provide a brief description of the aggregate productivity trends in the Netherlands.[6] Figure 1 shows the productivity growth pattern in the Netherlands, both for labor productivity and TFP. To smooth the productivity pattern we report five-year moving averages, and three-year moving averages to align with the firm-level analysis below. In addition, we consider the financial sector separately, mainly because our firm-level data used below contains information on non-financial firms only. Clearly, this sector shows a deviating pattern, with rising productivity figures up to the financial crisis, and a sharp drop afterwards. However, the total economy and the non-financial sector seem to be closely aligned. A slowdown of productivity growth can clearly be recognized for both TFP and labor productivity, although the most recent years seem to display a rebound. For direct comparison to the measures aggregated from the firm-level data that will be presented in later sections, we also show a three year moving average, starting in 2003. The slowdown is somewhat mitigated in these figures, but still both average labor productivity and TFP growth are 1 percentage point lower in the post-crisis years (2010-2016), as compared to the pre-crisis period (2001-2008).

[6] See also CPB (2017) and De Bondt (2015).

Figure 1. The productivity slowdown in the Netherlands

(a) Labor productivity (5 year moving average)

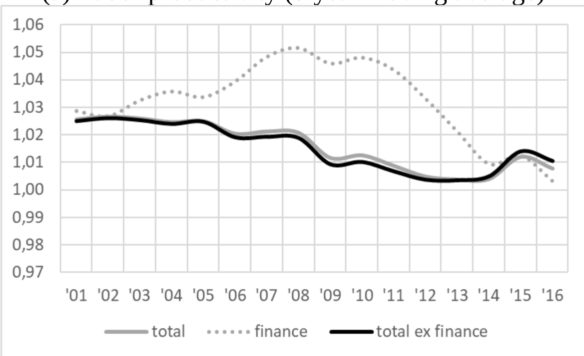

(b) TFP (5 year moving average)

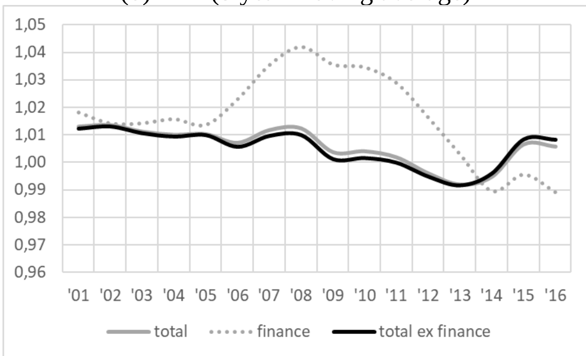

(c) Labor productivity (3 year moving average)

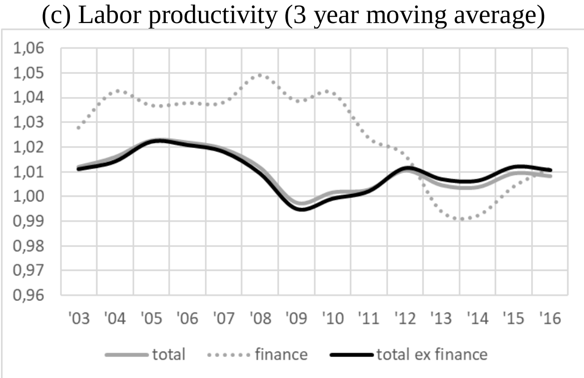

(d) TFP (3 year moving average)

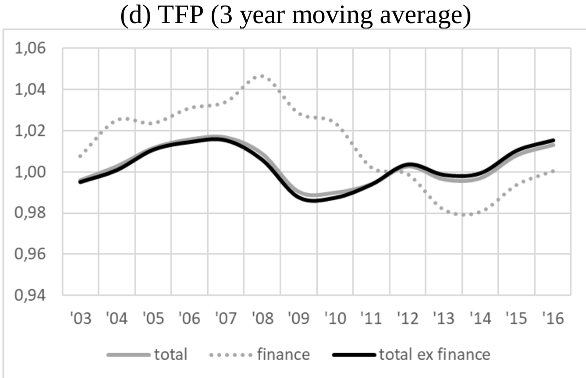

To assess whether the slowdown is a phenomenon across the whole economy, or can be attributed to specific industries, we next look at the industry-level figures of the Dutch growth accounts. Using the five-year moving averages, we determine the trend growth and trend-corrected average growth through a simple linear regression of labor productivity growth on an industry-specific constant and time-trend:

$$\Delta tfp_{jt} = \alpha_j + \beta_j t + \varepsilon_{jt}$$

Next, we compared the fitted (i.e. projected) growth of this regression for the estimation period ($\hat{\alpha}_j + \hat{\beta}_j t$)), with the counterfactual growth when trend growth would have been zero ($\beta_j = 0$) and annual growth would have been constant at the trend-corrected average growth ($\hat{\alpha}_j$). Figure 2 shows the difference in average annual growth under both scenarios during the sample period of 16 years across industries. Most industries have experienced a negative trend, with the agricultural sector being the exception with a mild positive trend. Overall the annual growth with the estimated trend is about 1 percentage point per year lower than without the estimated trend. In Information and communication, and to a lesser extent in Manufacturing, trend growth has been substantially negative, with about 4% and 2% lower growth as compared to the no-trend scenario.

Figure 2. Difference in average annual growth by industry with and without trend growth.

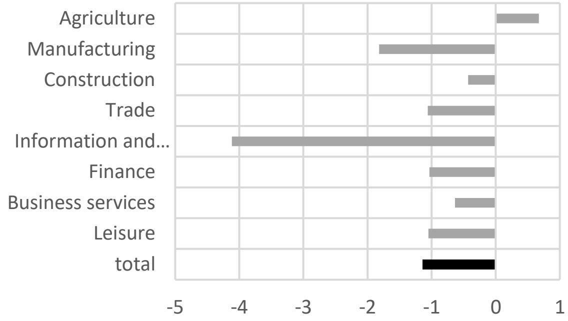

Figure 3. Productivity trends in the Netherlands against Europe and G7.

(a) Labor productivity, compared with EU28, Euro-area, and G7

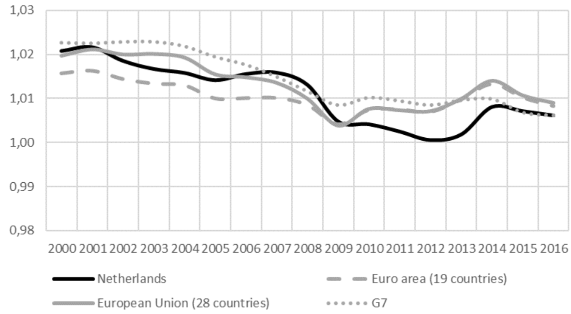

(b) Labor productivity compared with G7

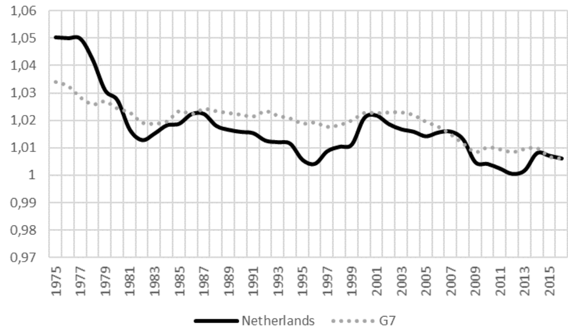

(c) TFP, compared to selected OECD countries

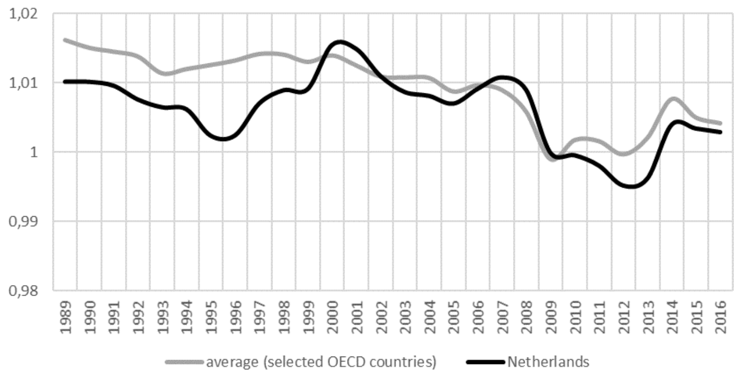

Source: OECD Productivity Database

Finally, we compare the productivity trend in the Netherlands to other countries. As already well documented, productivity growth has slowed down globally, and the European Union and the G7 economies are no exception. What is striking in figure 3a is that the Netherlands has experienced relatively high labor productivity growth before the crisis, but has fallen significantly under the international trend, only to catch up slightly during the last few years. Comparing to a longer time-series for the G7 economies, it is clear that the Netherlands has historically seen lower rates of labor productivity growth, with the exception of the seventies, when growth exceptionally high. Finally, comparing TFP growth to an average of selected OECD countries for which data is available, it can be seen that the Netherlands has been substantially below the OECD average during the nineties. After that TFP growth picked up and moved around the average, until the crisis, after which TFP growth has been relatively low again, compared to other OECD countries.

7. Firm-level heterogeneity in performance

7.1 Bottom-up calculation of the productivity slowdown

In this section we calculate aggregate productivity trends from the firm-level data. We start with firm-level productivity growth, thereby restricting the sample to continuing firms (i.e. the set of firms *C* in section 5). First, we determine the aggregate productivity growth for this group. Figure 4 shows different ways of calculating productivity growth from the firm-level data. The sample is restricted to continuing firms as to be able to calculate growth figures.[7] Besides the simple overall annual average, different kinds of weighting are employed. Size weights reflect the firm's size in overall (log) value added.[8] This gives larger firms more weight in the aggregate but does not control for any selection issues in the data. Sample weights do just this and weight the firm-level observations to reflect the share of firms in each size class/industry cell. This type of weighting therefore gives an average representative of the population. Finally, size and sample weighted figures combine both weights, and should therefore reflect the aggregate productivity, bearing in mind that the official aggregates are based on National Accounts and therefore subject to other corrections.

[7] In this version of the paper we do not go into different decomposition methods to calculate bottom-up productivity growth. Overall aggregate productivity growth is determined by growth of continuing firms, both in terms of productivity and relative size, as well as entry and exit.

[8] We experimented with different kinds of size weights. Ideally, for labor productivity we would use employment, as in the denominator of labor productivity, and for TFP we would use value added, not in logs. At this stage of our research however we were unable to get plausible results for the aggregate patterns. This is why log value added is used in the current version. The issue will be investigated in future versions.

Figure 4. Productivity growth (continuing firms, different kinds of weighting)

(a) Labour productivity (average)

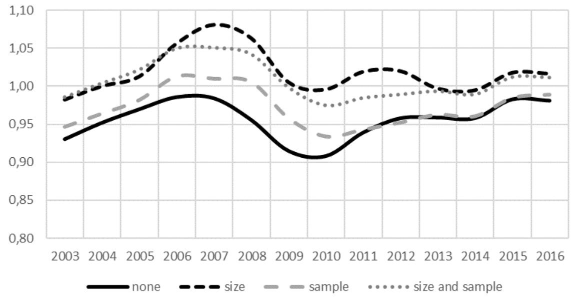

(b) Labour productivity (median)

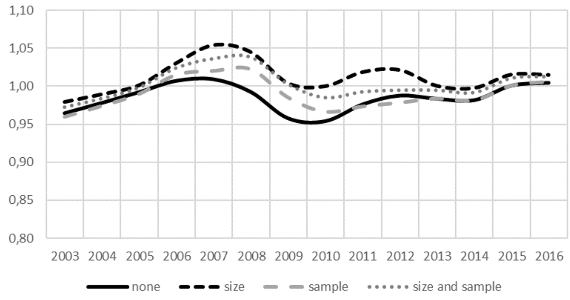

(c) TFP (average)

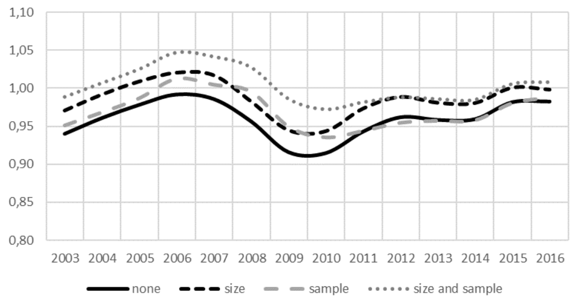

(d) TFP (median)

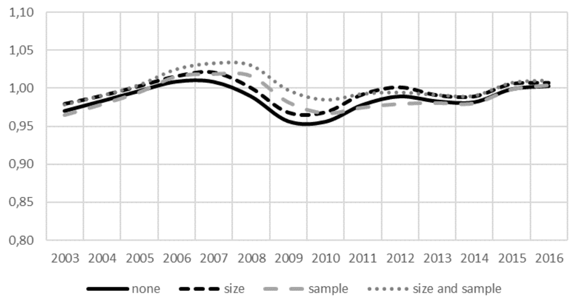

Overall, both the patterns for labour productivity and TFP growth seem to align with the aggregate pattern sketched above. Pre-crisis, the growth figures are higher and increasing, while at the time of the crisis productivity growth drops and remains stagnant afterwards. Several differences are worth to point however. In particular, the weighted series point stronger to a slowdown then the unweighted series. In the case of size weights, this suggests that especially larger firms have experienced this slowdown, and that this determines for a large part the aggregate slowdown. Naturally, the median exposes less variance than the average, while the labor productivity series is somewhat more volatile than TFP.

7.2 Productivity dispersion

Figure 5 shows the dispersion of TFP over time. To make the different statistics comparable, all series were normalized on the 2000 value.[9] Since the beginning of the century, dispersion has been gradually increasing up to the early crisis (2008), when dispersion dropped considerably, after which it initially started increasing again but then slowly seems to have stabilized. Panel 5.b focusses on the bottom

[9] Moreover, two changes in the definition of firm-level units were implemented by Statistics Netherlands (in 2006 and 2010) in the period considered, which seems to have affected the cross-sectional variance; we have filtered these effects by subtracting the average for each statistic in each of the pertinent periods.

and top decile. Here we can clearly see the reason for the drop in dispersion in the crisis and the stabilization in the more recent years. While TFP at the top decile moves rather smoothly over time, overall increases and decreases in dispersion seem to be driven by changes in the bottom decile. First, an increase in the dispersion in the first years of the century was caused by a sharp drop in relative productivity at the bottom decile, while the top decile only saw a slight decrease. Then, after a period of relatively high dispersion, the productivity of the lower decile caught up again to the level at the beginning of the century. Note that both the top and the bottom decile experience a decline in the levels of productivity, but that the lower decile is clearly more volatile. This can only be consistent with aggregate productivity growth (even if slowing down), when reallocation and/or entry/exit contributed positively to productivity growth. Another possibility is that average growth within the higher deciles increased, and by more than the average growth in the bottom decile.

Figure 5. Dispersion over time.

(a) Dispersion of TFP over time (2000 = 1)

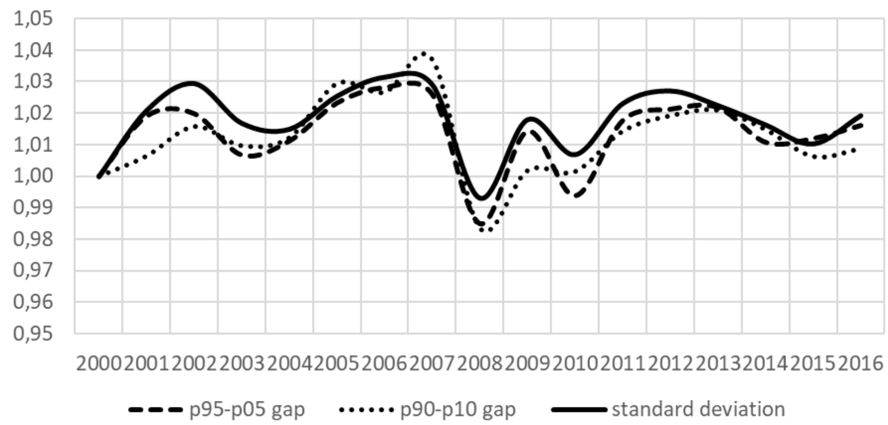

(b) Trends of bottom and top decile (2000 =1)

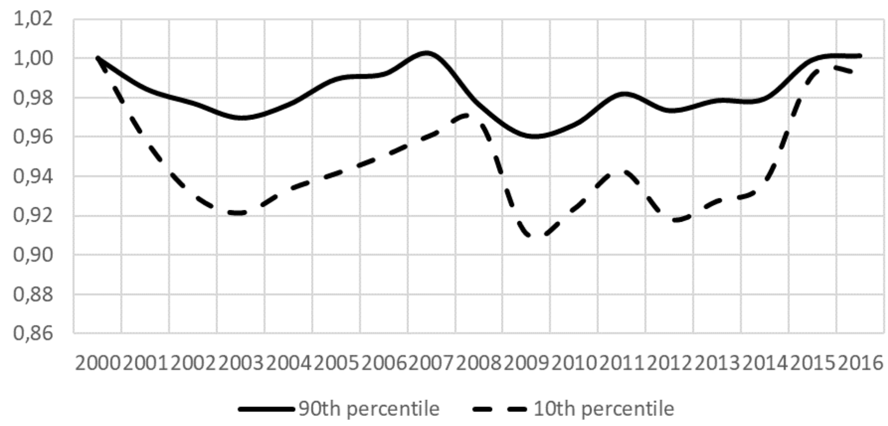

Figure 6. Current and forward TFP levels by decile.

(a) TFP growth by decile of the TFP distribution

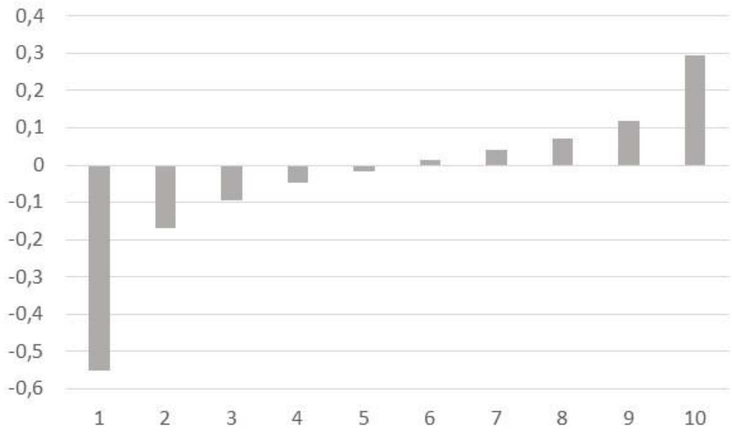

(b) Current and forward level of TFP by decile of the TFP distribution

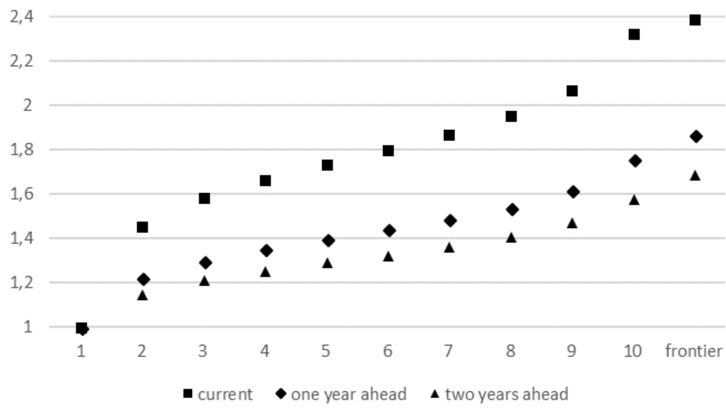

7.3 Persistence

An important question is whether the gap between firms at the frontier and those that are not is increasing. Using cross-country firm-level data from the ORBIS database, OECD (2015) finds that there is a divergence between the productivity of "the best versus the rest". In the previous section, we find that the spread of the productivity distribution in the Netherlands increased up to the crisis, then contracted as the crisis sets in, after which it increases again, and then narrows again towards the end of the period. In this sense, a divergence, in the sense of a widening productivity distribution, does not seem to be a continuous process in the Netherlands.

A widening distribution could be the consequence of a dynamic process in which laggard firms are being increasingly left behind. Figure 6 provides two pieces of evidence on this issue. The graph in panel (a) shows the TFP *growth* by decile of the TFP *level* distribution. This figure suggests that growth is highest at the frontier, and in fact contemporaneous productivity growth is a monotonously increasing function of the level, suggesting a widening of the distribution, with frontier firms moving further ahead of those behind.

The bottom panel, however, paints a different picture. The figure shows the TFP level by decile (normalized by the level in the bottom decile) in the current period, but also one and two years ahead *for the same firms* (i.e. the frontier is not re-determined in the next years). Frontier firms are about 239% more productive than those firms in the bottom decile in period *t*. While there is a high persistence of these firms remaining the most productive on average, the differences seems to narrow over time. After one year, the frontier firms of last year are on average 187% ahead of the bottom decile; after two years this is 169%. Moreover, a simple regression of productivity growth on the lagged and twice lagged TFP level shows that the correlation between the lagged levels of productivity and current growth rate is negative. This suggest that lagging firms, conditional on survival, are catching up, although a considerable distance remains. This result is consistent with OECD (2015) and CPB (2018), who report considerable dynamics up and off the frontier.

7.4 Firm profiles: high versus low productive firms

Firms at the frontier, do they do something special? Firms that move to the frontier, have they done something in particular to get there? In this section we look at different types of action and characteristics that may determine differences in productivity. Moreover, we look over time how these characteristics change over time across the productivity distribution. For this analysis we use the PS sample which is confined to firms with more 20 employees, and at a lower level of aggregation than the NFO sample (see section 4). Moreover, at this stage of writing the paper, the results reflect sample characteristics of the PS joint with the additional sources used to retrieve the firm characteristics, in particular the investment survey and the innovation survey. That is, there is no sample or size weighting. In future versions of the paper, the sensitivity of the results to weighting will be investigated. Finally, these data run up to 2014, but will be extended to include more recent year in following versions.

International activity

Increased globalization over the last decades has created both opportunities as well as threats for firms. In general, it is found that exporting firms have a higher productivity (Melitz, 2003) and have higher markups (De Loecker and Warzynski, 2012). The direction of causality may run both ways. On the one hand, it may be that only more productive firms can overcome a certain threshold to engage in exporting. Conversely, exporting firms may be exposed to higher competitive pressure and have access to a wider range of knowledge, stimulating productivity. Under both interpretations, a productivity slowdown would go hand in hand with a reduction of exporting activity. Table 1 presents some evidence on the difference between low and high productive firms in terms of productivity, and a change in exporting activities. The numbers confirm that exporting is more common in the higher deciles of the productivity distribution. Moreover, the difference with lower deciles in terms of the share of exporting firms has increased since the crisis. Looking back at figure 5, the increase in export active activity coincides with a period where the productivity distribution widened as the top decile firms kept up their productive performance, while firms at the lower end experienced a fall in productivity, suggesting a possible link of this divergence with an increased wedge in exporting activity.

Table 1. Differences in share of exporting firms between high and low productive firms.

exporter	lower	higher	Δ
2004-2008	0,48	0,51	0,04
2010-2014	0,58	0,66	0,08
Δ	0,10	0,15	
foreign owned	lower	higher	
2007-2010	0,06	0,10	0,05
2011-2014	0,07	0,15	0,08
Δ	0,02	0,05	

Another aspect of globalization is the rise of multinational companies. Foreign owned firms benefit from the experience and know-how of the controlling unit abroad, as well as of any sister affiliates. Bloom et al. (2011) for example show that U.S. owned firms perform better in terms of productivity due to their better use of ICT in their business processes. This is an example of cross-border knowledge flows through foreign affiliates.

The bottom panel of Table 5 confirms that higher productive firms have a higher probability of being foreign owned. Moreover, the share of multinational firms has increased over time in both groups, but particularly strong in the high productive group. Again, this suggest that part of the explanation of a divergence in productivity after the crisis could lie in the fact that foreign multinational firms increased in number, and kept up their productive performance, while domestic firms lost some ground, and their share was relatively higher in the bottom group.

Figure 6. Average firm age by decile of the productivity distribution

(a) Age by decile

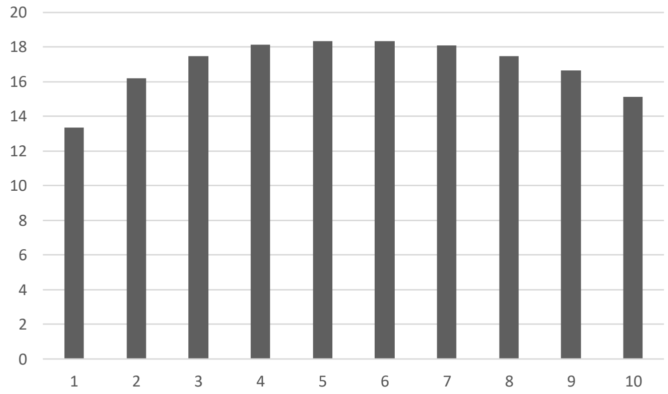

(b) Difference in average age between beginning (2000) and end of sample period (2014)

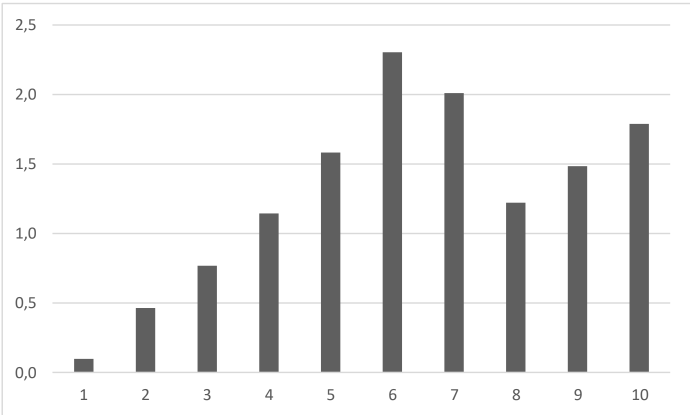

Firm age

Firm maturity is another aspect determining productivity. In the spirit of Schumpeterian creative destruction old firms are replaced by younger, more innovative and more productive firms. OECD (????) investigate start-up dynamics, and point out that an important source of aggregate growth is an economic environment where young successful start-up can thrive and grow out to larger employers. In this sense, a growing average age of firms could be a sign that entry and post-entry growth are being deterred, to the benefit of incumbents, ultimately hurting longer run growth.

Figure 6 shows the average age of firms along the productivity distribution. There is a clear inverted-U shape relation between age and productivity, indicating that firms become more productive as they grow older, but at some point this process of learning stops and reverses. Moreover, looking at the age of firms in each of the deciles, and comparing 2000 and 2014, figure 6 shows that over the sample period, higher productive firms have grown older on average. This could point to a lack of entry and post-entry success, holding back overall productivity growth.

Investment in intangibles and ICT

Finally, there is a growing and important literature on the importance of intangible investment (see Corrado et al. 2005, Haskel and Westlake, 2018, for an overview). Such investments include those beyond that in tangible assets, usually incorporating an element of knowledge and/or intellectual property. Examples include R&D, software, firm-specific human capital (i.e. training of workers), organizational capital (e.g. improvement in management practices), and brand equity (e.g. trademarks). The magnitude of this type of investment has been growing steadily over the last decades, and has surpassed investment in tangible investment in some developed countries.

Evidence for the importance of different types of innovation at the firm-level comes from Griliches (1979) and Crépon et al (1998) for R&D and innovation; Bresnahan et al. (2002) and Black and Lynch (2005) for organizational change. However, looking at table 2, which includes different types of innovation, and also investment in ICT hardware, the relation with productivity is not so clear-cut. Indeed, low and high productive firms do not seem to be different in these aspects. Three caveats need to be made in this respect. First and second, this analysis does not take into account the magnitude of the investments involved, and neither does it take into account any lagged effects in when investments become productive. Indeed, using a similar dataset as here, Mohnen et al. (2018) find that magnitude and timing matter, and we plan to extend the analysis in this direction. Third, it is possible that an effect of innovation on productivity can only be disentangles in a multivariate setting, controlling for other variables, as well as possible complementarities between innovation (see e.g. Polder et al. 2010 for evidence on this for the Netherlands).

Table 2. Differences in shares of firms by type of innovation and across periods.

		deciles				deciles	
R&D	1 to 5	6 to 10	Δ	process	1 to 5	6 to 10	Δ
2000-2006	0,18	0,17	-0,01	2000-2006	0,24	0,24	0,00
2008-2014	0,22	0,21	-0,01	2008-2014	0,32	0,30	-0,01
	0,04	0,05			0,08	0,06	
product	1 to 5	6 to 10		organizational	1 to 5	6 to 10	
2000-2006	0,28	0,29	0,01	2000-2006	0,33	0,31	-0,01
2008-2014	0,34	0,36	0,01	2008-2014	0,34	0,32	-0,02
Δ	0,06	0,06			0,01	0,01	
ICT	1 to 5	6 to 10					
2000-2006	0,40	0,35	-0,05				
2008-2014	0,42	0,34	-0,08				
Δ	0,02	-0,01					

8. Firm-level productivity and wages

More productive firms operate more cost efficiently than their competitors. How much of this cost efficiency translates into lower prices depends on the market structure under which a particular firm operates. In principle, firms with a productivity advantage can charge a higher markup, and be more profitable. This allows firms to pay their workers higher wages (rent sharing), although this also depends on the bargaining power of labor. Figure 7 shows two aspects of this. Firstly, average labor cost are clearly increasing in productivity. That is, more productive firms offer higher wages to their workers. The causality here is difficult to disentangle. Clearly, higher quality workers should increase a firm's productivity, and higher wages may be an indication of higher quality workers. Therefore, the causality here could run from wages to productivity, rather than the other way around.

Second, in terms of TFP, a firm in the top decile is on average about 3 times more productive than a firm in the bottom decile. For labor productivity, this about 2 times. Not the entire productivity differential translates into higher wages however, since labor cost per worker are about 1.6 times higher in the top decile. This is also confirmed by the fact that the average productivities increase more when moving up the distribution, than the labor cost do.

Figure 7. Average labor productivity, TFP, and labor cost per worker by decile of the TFP distribution.

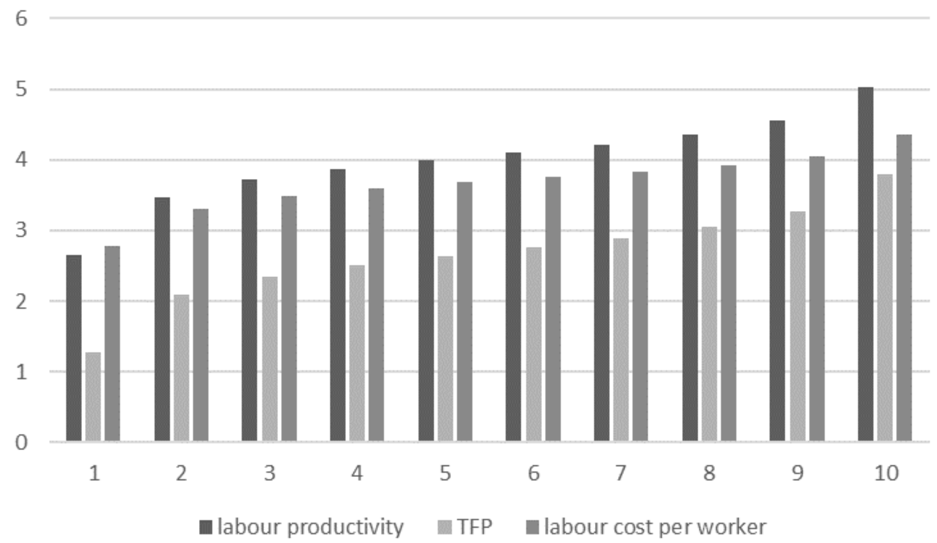

Figure 8. Wage differential between bottom and top decile over time.

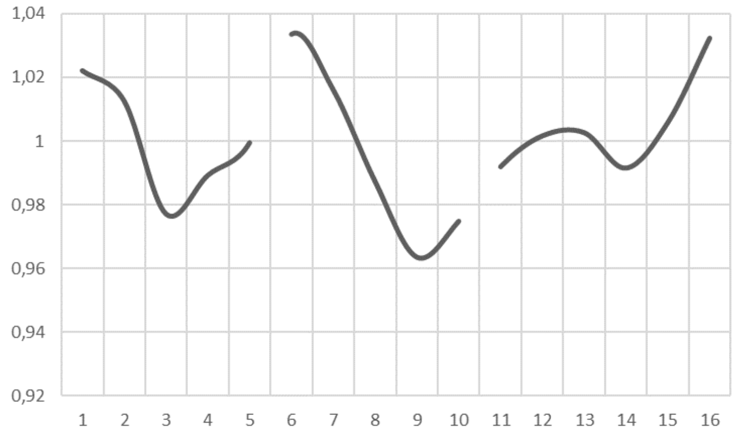

Note: The years 2005-2006 and 2010-2011 are not connected, because of break in the data regarding the reporting units in the firm-level data.

Figure 8 presents evidence that the wage differential has been increasing since the crisis, especially in 2015 and 2016. This matches the period of the macro-economic productivity upturn (figure 1), and suggests that especially workers in firms at the top end of the productivity distribution have benefited from this upturn. The widening wage differential follows a periods of narrowing differences, from 2006 up to 2009.

9. Summary of preliminary findings and to do list

This work-in-progress paper looks at the firm-level distribution of productivity in the Netherlands in the light of the macro-economic productivity slowdown. In this section we summarize some of the (preliminary) results.

We document the heterogeneity in productivity performance surrounding the macro and industry-level figures from around the beginning of the century. We find that the aggregate productivity trend as described by the National Accounts data can be mimicked by a bottom-up calculation using firm-level data, and taking into account firm size and sample weights, with larger firms having a bigger role in the slowdown. Moreover, productivity levels at the top of the distribution seem to have been rather stable, while at the bottom of the distribution there is significantly more volatility, so that changes in the shape of productivity distribution related to the lower end of the distribution seem to have contributed more to the slowdown. This latter finding goes hand in hand with a pattern of contraction and widening of the productivity distribution over time.

We also find that while productivity growth increases with the level of productivity, the differences between low and high productivity firms one and two years ahead decreases. Therefore, while there is a high persistence at the frontier, there is no divergence over time, conditional on survival.

Next, we looked at characteristics of firms along the productivity distribution. Highly productivity firms are more likely to export and to be part of a multinational company. The share of internationally active firms seems to have increased in the higher end of the productivity distribution, but not in lower end, thereby contributing to a divergence in productivity after the crisis.

The relation of productivity with age is non-linear. Firms at the frontier have grown older on average during the period under consideration, possibly signaling a lack of entry and post-entry growth, which could also be holding back overall productivity growth. Finally, we do not find clear-cut differences in investment in intangible assets or ICT. Taking into account the timing and magnitude of these investments is on the agenda for further research.

Finally, we document that high productive firms pay higher wages, although the productivity differential between high and low productive firms does not seem to translate fully into wage differential, suggesting that rent sharing is only partial. However, this wage differential has been increasing over time. suggesting that especially workers in firms at the top end of the productivity distribution have benefited from the recent upturn in productivity growth in the Netherlands.

As mentioned, this paper is work in progress. The analysis will be refined and extended in various directions. The main concern is to complement the bottom-up approach with entry and exit to get a more complete picture of the business dynamics underlying the aggregate pattern, as described by the simple conceptual in section 2. Moreover, in this bottom-up approach it is possible to distinguish the contribution of different firms, classified along the characteristics analyzed in section 7. Besides updating the data in this section with more recent recent years, these characteristics can be extended and refined, for example using information on imports and foreign affiliates, different types of ICT, financial information, worker educational attainment, introducing lags, and taking into account magnitudes of investment.

References (incomplete and to be updated)

Bartelsman and Doms, 2000, Understanding Productivity: Lessons from Longitudinal Microdata, Journal of Economic Literature, 38(3): 569-594.

Byrne, Fernald, and Reinsdorf, 2016, Does the United States have a productivity slowdown or a measurement problem?, BPEA Conference Draft, March 10-11, 2016.

Cindy Cunningham & Lucia Foster & Cheryl Grim & John Haltiwanger & Sabrina Wulff Pabilonia & Jay Stewart & Zoltan Wolf, 2018. "Dispersion in Dispersion: Measuring Establishment-Level Differences in Productivity," Working Papers 18-25, Center for Economic Studies, U.S. Census Bureau.

Decker et al. 2018

Gordon, 2016, The Rise and fall of American Growth: The U.S. Standard of Living Since the Civil War.

Haan, Mark de, Erik Veldhuizen, Murat Tanriseven and Myriam van Rooijen-Horsten (2014), “the Dutch growth accounts: measuring productivity with non-zero profits”, Review of Income and Wealth, Series 60, Supplement Issue.

Mortensen, 2003, Wage Dispersion: Why are Similar Workers Paid Differently?

OECD, 2015 , The future of productivity.

OECD, 2016, The Best versus the Rest: The Global Productivity Slowdown, Divergence across Firms and the Role of Public Policy.

OECD, 2017a, OECD Compendium of Productivity Indicators 2017.

OECD, 2017b, The great divergence(s).

ONS (2017)

Summers, 2016, The age of secular stagnation, Foreign affairs.

Syverson, 2011, What determines productivity growth?, Journal of Economic Literature.