

Ilmenau Economics Discussion Papers, Vol. 28, No. 180

**The Fallacy in Productivity Decomposition
(Updated June 2023)**

Simon Bruhn, Thomas Grebel, Lionel Nesta

June 2023

Institute of Economics

Ehrenbergstraße 29
Ernst-Abbe-Zentrum

D-98 684 Ilmenau

Phone 03677/69-4030/-4032

Fax 03677/69-4203

<https://www.tu-ilmenau.de/iedp>

ISSN 0949-3859

This preprint has not undergone peer review or any post-submission improvements or corrections. The Version of Record of this article is published in the Journal of Evolutionary Economics, and is available online at <https://doi.org/10.1007/s00191-023-00816-8>

The Fallacy in Productivity Decomposition

(Updated June 2023)

Simon Bruhn* Thomas Grebel† Lionel Nesta‡

This paper argues that the typical practice of performing growth decompositions based on log-transformed productivity values induces fallacious conclusions: using logs may lead to an inaccurate aggregate growth rate, an inaccurate description of the microsources of aggregate growth, or both. We identify the mathematical sources of this log-induced fallacy in decomposition and analytically demonstrate the questionable reliability of log results. Using firm-level data from the French manufacturing sector during the 2009-2018 period, we empirically show that the magnitude of the log-induced distortions is substantial. Depending on the definition of accurate log measures, we find that around 60-80% of four-digit industry results are prone to mismeasurement. We further find significant correlations of this mismeasurement with commonly deployed industry characteristics, indicating, among other things, that less competitive industries are more prone to log distortions. Evidently, these correlations also affect the validity of studies that investigate the role of industry characteristics in productivity growth.

Keywords: productivity decomposition, growth, log approximation, geometric mean, arithmetic mean

*Researcher, Chair for Economic Policy, Institute of Economics, Ilmenau University of Technology, Germany, email: simon.bruhn@tu-ilmenau.de

†Professor, Chair for Economic Policy, Institute of Economics, Ilmenau University of Technology

‡Professor, GREDEG - CNRS - University Côte d'Azur, OFCE SciencesPo and SKEMA Business School, 250 rue Albert Einstein – CS 10269, Sophia Antipolis 06905, Cedex, France

1. Introduction

This paper questions the typical practice of performing growth decompositions based on the log-transformed values of productivity. We argue that representing firm-level productivity in logs may lead to an inaccurate aggregate growth rate, an inaccurate description of the microsources of aggregate growth, or both. These three cases of potential misconceptions are what we refer to as the fallacy in productivity decomposition. Therefore, policy recommendations stemming from log-based decomposition exercises may prove inappropriate.

Productivity decomposition methods are useful tools to shed light on the underlying causes of aggregate productivity movements. The most commonly used shift-share decomposition methods include those proposed by Griliches and Regev (1995), Foster, Haltiwanger and Krizan (2001) and Melitz and Polanec (2015). Whereas the former two are time-series approaches based on the seminal contribution by Baily, Hulten and Campbell (1992), the latter is based on the cross-sectional methodology by Olley and Pakes (1996). Despite their technical differences, they all use the weighted average of firm-level productivity and decompose aggregate productivity growth according to its underlying microsources; for firm-level analyses, these include (i) productivity changes at the individual firm level (within-firm effect), (ii) shifts in market shares between firms (between-firm effect)¹, (iii) entries of new firms, and (iv) exits of incumbents.

The use of these methods differs in various ways. Some studies use labor productivity, while others use total factor productivity; some use inputs, while others choose output shares as weights (see e.g., Fagerberg, 2000; Foster et al., 2001; Melitz and Polanec, 2015; Decker et al., 2017). Discrepancies may also arise due to the chosen length of the period analyzed, as a decomposition of shorter periods typically yields larger within-firm contributions (Brown et al., 2018). A further methodological difference is whether to measure firm-level productivity in levels or in logs, with the typical practice being the use of logs (see e.g., Van Biesebroeck, 2008; Melitz and Polanec, 2015).² However, as we argue in this paper, representing firm-level productivity in logs entails the risk of severe misinterpretations regarding aggregate productivity growth as well as its decomposed elements. We identify three underlying sources: (i) the log approximation error, as a consequence of the logarithm's concavity; (ii) the reference deviation, arising from a different reference assumption implicit in log differences; and (iii), the mean deviation, caused by the difference in the deployed benchmark productivity.

In our analysis, we follow the standard textbook definition of aggregate productivity

¹Foster et al. (2001) apply a third component for incumbent firms, which consists of an interaction between the first two components (cross-firm effect).

²According to our literature review, no less than two-thirds of the identified firm-level studies use logs when decomposing labor productivity. See Table A1.

growth which reads as follows

$$\hat{\Phi} = \frac{\Phi_2 - \Phi_1}{\Phi_1}, \quad (1)$$

where Φ_1 and Φ_2 denote aggregate productivity in two successive time periods. Aggregate productivity is simply the sum of firm-level output over the sum of firm-level input(s). This definition is the most conventional measure of aggregate productivity, and we take $\hat{\Phi}$ as being the most accurate measure of aggregate productivity growth. As mentioned, however, many studies choose to measure firm-level productivity in logs which inevitably implies a deviation from the above definition.

At the outset, we want to emphasize that the composition of an aggregate can naturally be freely defined by the researcher. Hence, when aggregate productivity is defined as a share-weighted aggregate of log firm-level productivity, the mathematics are certainly consistent. However, from an economic viewpoint, productivity is essentially a measure of efficiency, that is, the relationship between output and input(s). Even though the log of *firm-level* output over inputs might still be a reasonable measure of *firm-level* efficiency, the use of a log *aggregate* for an industry’s or a country’s *aggregate* efficiency is questionable, as it lacks a clear link to the relationship between aggregate output and input(s) (see e.g., Melitz and Polanec, 2015). Moreover, productivity components based on log measures, even if reported in log points, are frequently interpreted as percent changes in the literature.³ This reinforces our perception that logs are in fact considered an approximation to the representation of growth rates in percent, that is, based on levels.

In our argument, we aim in a similar direction as Dias and Marques (2021a) who shed light on the potential misconceptions in productivity decompositions caused by the fact that the aggregation of logs leads to a geometric instead of an arithmetic mean.⁴ With respect to their contribution, we believe our work to contribute in several ways. First, instead of the aggregate perspective, we analyze the impact of logs from the perspective of an individual firm’s contribution to aggregate growth. Adopting the firm-level perspective reveals that the discrepancy between level and log results for individual productivity components is not only caused by the different types of means but that it can be traced back to three sources of distortions mentioned above, namely, the log approximation error, the reference deviation, and the mean deviation. The separation of the three log distortions provides a straightforward analytical framework to determine the circumstances under which firm-level contributions are over- or underestimated by logs. Second, we extend the findings by Dias and Marques, who use a modified version of the Melitz and Polanec (2015) decomposition method, by setting out the log distortions in

³See Van Biesebroeck (2005); Foster et al. (2006); Haskel and Sadun (2009) or Melitz and Polanec (2015) who substantiate the widespread use of logs to approximate productivity growth rates; the latter, for instance, explicitly remarks that “... [a]ll productivity changes are reported as log percents (or log points) – and can thus be interpreted as percentage point changes” (Melitz and Polanec, 2015, p. 371).

⁴The fact that logs lead to a geometric mean has previously also been pointed out by, for example, Van Biesebroeck (2008) or Brown et al. (2018).

the widely applied decomposition methods by Foster et al. (2001) (henceforth: FHK). We further show that our findings also hold for the method by Griliches and Regev (1995) and the decomposition (in its original form) by Melitz and Polanec (2015). Third, we show how the log distortions differ between studies conducted for the average industry and for individual industries, revealing that the average industry is less affected by log distortions. As most studies typically investigate the average industry, these findings may be of particular interest. Fourth, we quantify the fallacy in decomposition by exploiting our data on four-digit industries and thereby provide striking evidence for the scope of our findings. Fifth, we document how log distortions relate to certain industry characteristics, revealing, in particular, that the lack of competition and a high degree of industry openness reinforce log distortions. We believe we thereby take an important step to evaluate the extent to which past and future studies of industry dynamics may be affected by the use of logs and how they compare.

The remainder of the paper is structured as follows. In Section 2, we define and formalize the first two distortions, namely, the log approximation error and the reference deviation. In Section 3, we discuss the use of logs in the productivity components of the FHK decomposition method, which reveals the mean deviation as the third log-induced distortion. Section 4 shows the magnitude of the log distortions using firm-level data of the French manufacturing sector. In Section 5, we quantify the fallacy in decomposition and address how the log-induced discrepancies correlate with certain industry characteristics. Section 6 concludes.

2. Log approximation error and reference deviation

It is common practice in productivity decompositions to represent firm-level productivity in logs. The main motive lies in the linearization of the decomposition exercise (Van Biesebroeck, 2008), achieved by using a log difference as an approximation to a productivity growth rate:

$$\frac{\varphi_{i2} - \varphi_{i1}}{\varphi_{i1}} \approx \ln(\varphi_{i2}) - \ln(\varphi_{i1}) \quad (2)$$

where φ_{i1} and φ_{i2} denote productivity levels of firm i in two successive periods. Due to the concavity of the logarithmic function, logged values underestimate productivity growth, i.e., $(\varphi_{i2} - \varphi_{i1})/\varphi_{i1} - (\ln(\varphi_{i2}) - \ln(\varphi_{i1})) \geq 0$. It is this well-known deviation between the level growth rate in the form of a ratio and the log difference, which we refer to as the **log approximation error**.

Proposition 1 The use of logs introduces a log approximation error, that is, a systematic underestimation of productivity growth rates.

It is usually argued that the occurring approximation error can be kept within reasonable limits as long as the growth rates fluctuate within a range of approximately $\pm 10\%$. However, it is not uncommon for individual firms to experience a change in productivity beyond such values. Moreover, while the positive and negative growth rates of firms may partly balance out in the aggregate, the log approximation errors will not because they are consistently positive, independent of whether a firm increases or decreases its productivity. In addition, because each decomposition component reflects a weighted sum and involves a different set of weights, the implied log-approximation error is scaled accordingly, thus making predictions about the magnitude of the error illusive.

Aside from the log approximation error, there is a further discrepancy between levels and logs which appears when aggregating firm-level growth rates. In contrast to levels where absolute changes in firm-level productivity are measured against some reference productivity – usually the aggregate productivity of the previous period, as shown in eq. (1) – the reference of a productivity growth rate calculated with log-transformed values is implicit in the firm-individual log difference. Hence, we have to address what we call the **reference deviation**. For each firm-specific productivity growth rate, the reference productivity will differ from the unique reference productivity applied in the case of levels.

Proposition 2 Aggregating log differences as a proxy for growth rates induces a reference deviation arising from the idiosyncratic reference productivity when using logs in contrast to a single reference productivity when calculating growth rates in levels.

To illustrate the two so far mentioned log distortions, we start by revisiting the definition we gave in the introduction, defining aggregate productivity as the ratio of aggregate output to aggregate input(s). In the case of aggregate labor productivity, this reads as:⁵

$$\Phi_t = \frac{\sum Y_{it}}{\sum L_{it}} \quad (3)$$

where Y_{it} denotes output (e.g., value-added) and L_{it} input (e.g., working hours) of firm i at time t . Following the decomposition method by Foster et al. (2001), aggregate productivity is calculated as the share-weighted mean of firm-level productivity; calculated in levels, this implies:

$$\hat{\Phi}_{lev} = \left(\sum s_{i2} \cdot \varphi_{i2} - \sum s_{i1} \cdot \varphi_{i1} \right) \frac{1}{\Phi_1} \quad (4)$$

where φ_{i1} and φ_{i2} denote productivity levels of firm i in two successive periods, and s_{i1} and s_{i2} indicate share weights. To ensure that the aggregation of firm-level data corresponds

⁵For the sake of simplicity, we confine our investigation to labor productivity as a productivity measure and will not address total factor productivity (TFP) because its measurement poses additional challenges (Dosi et al., 2015; Dosi and Grazzi, 2006), although it does not hurt the general statements in this paper regarding the causes of log distortions.

to the industry aggregate as defined in eq. (3), we use input shares for weighing individual firm productivity, i.e., $s_{it} = \frac{L_{it}}{\sum L_{jt}}$.⁶ Note that each firm's absolute productivity growth is divided by the reference productivity Φ_1 , which renders the aggregate productivity growth rate $\hat{\Phi}_{lev}$.⁷

If, instead, firm-level productivity is measured in logs, the share-weighted industry aggregate is defined as $\Phi_{t,log} = \sum s_{it} \cdot \ln \varphi_{it}$ (see e.g., Van Biesebroeck, 2008; Bartelsman et al., 2013; Melitz and Polanec, 2015; Decker et al., 2017). As the difference between two log aggregates corresponds to a percentage change, aggregate productivity growth in logs can be expressed as follows:

$$\hat{\Phi}_{log} = \sum s_{i2} \cdot \ln \varphi_{i2} - \sum s_{i1} \cdot \ln \varphi_{i1} \quad (5)$$

which is equivalent to the log difference of two geometric means:

$$\hat{\Phi}_{log} = \ln \prod \varphi_{i2}^{s_{i2}} - \ln \prod \varphi_{i1}^{s_{i1}} \quad (6)$$

In other words, instead of a growth rate between two share-weighted arithmetic means, logs approximate a growth rate between two share-weighted geometric means (Van Biesebroeck, 2008; Brown et al., 2018; Dias and Marques, 2021a). This approach inevitably affects the computation of productivity growth because a geometric mean is more sensitive to smaller numbers than to larger numbers, which mitigates the effect of high values while reinforcing the impact of low values. Moreover, as Jensen's inequality proves, the (weighted) geometric mean is always smaller than the (weighted) arithmetic mean unless all numbers constituting the means are equal (Casella and Berger, 2002).

However, this does not imply that the growth rate between two arithmetic means must also be larger, as this is contingent on the composition of the respective means and their changes over time, as shown by Dias and Marques (2021a). Put differently, it depends on the development of each firm's productivity and input share while considering the individual firm's positioning within the industry's initial productivity distribution. To further elaborate on this point, we adopt the perspective of the individual firm. From equations (4) and (5), we can extract the individual firm productivity contribution C_i to aggregate productivity growth measured in levels and logs, respectively:

$$C_{i,lev} = (s_{i2}\varphi_{i2} - s_{i1}\varphi_{i1}) \cdot \frac{1}{\Phi_1} \quad (7)$$

$$C_{i,log} = s_{i2} \cdot \ln \varphi_{i2} - s_{i1} \cdot \ln \varphi_{i1} \quad (8)$$

⁶In using input shares, we follow the *denominator rule in share-weighting aggregation* of Färe and Karagiannis (2017), who show that, when aggregating, consistent results are achieved only by using the denominator of the productivity measure as weights.

⁷In Appendix D, we indicate how our findings apply when the time average productivity, $\bar{\Phi}$, is used as a reference, as, for instance, that by Melitz and Polanec (2015).

The comparison of $C_{i,lev}$ and $C_{i,log}$ reveals how the computation of firm-level contributions to aggregate growth is affected by the two log distortions we identified in propositions (1) and (2). The log approximation error is visible in the ratio in eq. (7) as opposed to the log difference in eq. (8), while the reference deviation shows in the different reference productivities, which is the aggregate productivity, Φ_1 , in eq. (7) and the firm's initial productivity, φ_{i1} , in eq. (8).

Hence, as mentioned above, a log difference embodies its individual reference productivity in the respective subtrahend. In eq. (8), the subtrahend is $\ln(\varphi_{i1})$, implying that the log difference uses each firm's initial productivity φ_{i1} as a reference productivity.⁸ Consequently, individual firms' productivity growth rates measured in logs can have both a higher or a lower contribution to aggregate growth compared to the level results, depending on the positioning of each firms' initial productivity, φ_{i1} , relative to aggregate productivity, Φ_1 . Therefore, the reference deviation plays a crucial role in determining whether a firm's contribution to aggregate growth is underestimated, i.e., $C_{i,lev} > C_{i,log}$, or overestimated, i.e., $C_{i,lev} < C_{i,log}$, by logs.⁹

Following this line of reasoning, when aggregating firm contributions to aggregate productivity growth based on logs, the resulting aggregate growth rate is susceptible to bias. It is the industry structure and its change over time that determines the extent to which productivity growth contributions are under- or overestimated, which raises doubts about the reliability and comparability of productivity measures based on logs.

In the following section, we flesh out our propositions by decomposing aggregate productivity growth according to the FHK method. For each productivity component in the FHK method, we propose a separation of the different log distortions, which provides a straightforward analytical approach to determine the circumstances under which firms' productivity contributions are over- or underestimated by logs.

3. Decomposing the log distortions in the FHK decomposition

The FHK decomposition distinguishes three groups of firms, i.e., surviving (S), entering (N), and exiting firms (X). The contribution of surviving firms is further broken down into three subcomponents, which they label the within-firm effect (WFE), the between-

⁸Note that we draw this analogy based on the approach that a log difference between φ_{i2} and φ_{i1} *approximates* the ratio of the absolute difference ($\varphi_{i2} - \varphi_{i1}$) and the *initial productivity* φ_{i1} , while we isolate the inaccuracy caused by the log approximation in the approximation error. Hence, our analogy is not opposed to Törnqvist et al. (1985), who stated that a log difference *equals* the ratio of the absolute difference ($\varphi_{i2} - \varphi_{i1}$) and the *logarithmic mean* $L(\varphi_{i1}, \varphi_{i2})$, with $L(\varphi_{i1}, \varphi_{i2}) = (\varphi_{i2} - \varphi_{i1}) / \ln(\varphi_{i2}/\varphi_{i1})$ and $(\varphi_{i1}\varphi_{i2})^{1/2} < L(\varphi_{i1}, \varphi_{i2}) < (\varphi_{i1} + \varphi_{i2})/2$ for $\varphi_{i1} \neq \varphi_{i2}$.

⁹To illustrate, let us assume two firms with constant market shares $s_{it} = 10\%$ for $i=\{1,2\}$, productivity levels $\varphi_{11} = 50$ and $\varphi_{21} = 150$, and an initial aggregate productivity $\Phi_1 = 100$. Suppose both firms increase their productivity by 10%. With logs, the contribution of both firms to aggregate productivity growth will be the same, namely, $0.1 \cdot \ln(1.1) \approx 0.953\%$. Calculated in levels, the impact of firm 1 will be smaller than the impact of firm 2, namely, $0.1 \cdot \frac{55-50}{100} = 0.5\%$ for firm 1 and $0.1 \cdot \frac{165-150}{100} = 1.5\%$ for firm 2. Hence, logs overestimate the impact of firm 1 and underestimate the impact of firm 2.

firm effect (BFE), and the cross-firm effect (CFE). Expressed in levels, the decomposition reads as follows:

$$\begin{aligned}
\hat{\Phi}_{lev} &= \sum_{i \in S} \underbrace{\frac{1}{\Phi_1} \cdot s_{i1} \cdot \Delta \varphi_i}_{WFE_{i,lev}} + \sum_{i \in S} \underbrace{\frac{1}{\Phi_1} \cdot \Delta s_i \cdot (\varphi_{i1} - \Phi_1)}_{BFE_{i,lev}} \\
&+ \sum_{i \in S} \underbrace{\frac{1}{\Phi_1} \cdot \Delta s_i \cdot \Delta \varphi_i}_{CFE_{i,lev}} + \sum_{i \in N} \underbrace{\frac{1}{\Phi_1} \cdot s_{i2} \cdot (\varphi_{i2} - \Phi_1)}_{N_{i,lev}} \\
&+ \sum_{i \in X} \underbrace{\frac{1}{\Phi_1} \cdot s_{i1} \cdot (\Phi_1 - \varphi_{i1})}_{X_{i,lev}}
\end{aligned} \tag{9}$$

In logs, the individual productivity components are expressed as follows:

$$\begin{aligned}
\hat{\Phi}_{log} &= \sum_{i \in S} \underbrace{s_{i1} \cdot \Delta \ln \varphi_i}_{WFE_{i,log}} + \sum_{i \in S} \underbrace{\Delta s_i \cdot (\ln \varphi_{i1} - \Phi_{1,log})}_{BFE_{i,log}} \\
&+ \sum_{i \in S} \underbrace{\Delta s_i \cdot \Delta \ln \varphi_i}_{CFE_{i,log}} + \sum_{i \in N} \underbrace{s_{i2} \cdot (\ln \varphi_{i2} - \Phi_{1,log})}_{N_{i,log}} \\
&+ \sum_{i \in X} \underbrace{s_{i1} \cdot (\Phi_{1,log} - \ln \varphi_{i1})}_{X_{i,log}}
\end{aligned} \tag{10}$$

Note that in the BFE component as well as in the components of entering and exiting firms, firm-level productivity is set in relation to aggregate productivity as a benchmark, which is Φ_1 for levels and $\Phi_{1,log}$ for logs. As noted in Section 2, we can rewrite the log aggregate as $\Phi_{1,log} = \sum s_{i1} \ln \varphi_{i1} = \ln \prod \varphi_{i1}^{s_{i1}}$. We will denote the geometric mean of firm-level productivity as Π_1 , i.e., $\Phi_{1,log} = \ln \prod \varphi_{i1}^{s_{i1}} = \ln \Pi_1$, distinguishing it from the arithmetic mean used by levels, Φ_1 . Self-evidently, the discrepancy in means will induce a further distortion in the computation of productivity growth, which we term **mean deviation**.

Proposition 3 Logs introduce a mean deviation in the between-firm effect and in the components of entering and exiting firms, as logs use a geometric mean instead of an arithmetic mean as the benchmark productivity.

As this proposition emphasizes, the mean deviation will only surface in the BFE and in the components of entering and exiting firms. The impact of the mean deviation on the discrepancy between the log and the level *aggregate* growth rate, however, will be zero because market shares sum to one in each period, i.e., $\sum s_{i1} = \sum s_{i2} = 1$ (see e.g., Melitz and Polanec, 2015). Additionally, note that in a balanced panel, the impact of the mean deviation on the *aggregate* between-firm effect is always zero (see Section 3.2) since the

sum of changes in shares is zero, i.e., $\sum_{i \in S} \Delta s_i = 0$ (see e.g., Baily et al., 2001).

The following sections set out the extent to which the three identified log distortions affect the individual components within the FHK decomposition method, by which we extend the analysis provided by Dias and Marques (2021a). While sharing similarities in the decomposition of the log distortions in the WFE and the CFE, we also provide a decomposition for the three remaining components.

3.1. Log distortions in the within-firm effect

The within-firm effect (WFE) depicts the part of aggregate growth driven by productivity improvements at the individual firm level, weighted by each firm's initial input share. Comparing the WFE from equations (9) and (10), and building upon our findings in Section 2, the level and log results concerning the WFE diverge as a consequence of a combination between the log approximation error and the reference deviation:

$$\begin{aligned}
\varepsilon_{i,W} &= WFE_{i,lev} - WFE_{i,log} \\
&= s_{i1} \left(\frac{\Delta \varphi_i}{\Phi_1} - \Delta \ln \varphi_i \right) \\
&= \underbrace{s_{i1} \left(\frac{\Delta \varphi_i}{\varphi_{i1}} - \Delta \ln \varphi_i \right)}_{\varepsilon_{i,W,appr}} + \underbrace{s_{i1} \left(\frac{\Delta \varphi_i}{\Phi_1} - \frac{\Delta \varphi_i}{\varphi_{i1}} \right)}_{\varepsilon_{i,W,ref}}
\end{aligned} \tag{11}$$

As shown in eq. (11), a further decomposition of the difference between the WFE in levels and logs separates the approximation error ($\varepsilon_{i,W,appr}$) from the reference deviation ($\varepsilon_{i,W,ref}$) – the latter arising from the alternative use in reference productivities.

Since changes in shares are ignored in the WFE, the input share s_{i1} simply works as a scaling factor of the distortions. The log approximation error ($\varepsilon_{i,W,appr}$) is increasing in $|\Delta \varphi_i|$. Due to the concavity of the logarithm and the fact that input shares must always be positive, $s_{i1} > 0$, the approximation error $\varepsilon_{i,W,appr}$ is always nonnegative ($\varepsilon_{i,W,appr} \geq 0$). Hence, the log approximation error introduces a systematic underestimation in the WFE component. The sign and magnitude of the reference deviation ($\varepsilon_{i,W,ref}$) depend on the position of the firm's productivity and its development within the industry's productivity distribution, i.e., on the relationship of φ_{i1} and Φ_1 and on the directional change in $\Delta \varphi_i$. $\varepsilon_{i,W,ref}$ is positive, if $(\varphi_{i1} > \Phi_1 \wedge \Delta \varphi_i > 0) \vee (\varphi_{i1} < \Phi_1 \wedge \Delta \varphi_i < 0)$, and negative if $(\varphi_{i1} > \Phi_1 \wedge \Delta \varphi_i < 0) \vee (\varphi_{i1} < \Phi_1 \wedge \Delta \varphi_i > 0)$.

Summarizing the log distortions in the WFE, the reference deviation will add to the consistently positive log approximation error, compensate or even overcompensate the error, if $\varepsilon_{i,W,appr} + \varepsilon_{i,W,ref} < 0$. As noted by Dias and Marques (2021a), the described

tendencies in the two error terms anticipate a generally positive log distortion in the WFE. This is also reflected in our empirical findings (Section 4).

3.2. Log distortions in the between-firm effect

The between-firm effect (BFE) depicts input share fluctuations for the analysis of reallocation effects within the industry. Although the BFE holds the productivity measure constant focusing on changes in shares, we can make use of the fact that, from a mathematical perspective, normalizing the impact of the BFE using the weighted arithmetic mean Φ_1 and the weighted geometric mean Π_1 , respectively, corresponds to the mathematical equivalent of a productivity growth rate (see equations 9 and 10). Due to this analogy of a growth rate, we can decompose the log distortions in the BFE according to propositions (1) to (3):

$$\begin{aligned}
\varepsilon_{i,B} &= BFE_{i,lev} - BFE_{i,log} \\
&= \Delta s_i \left(\frac{\varphi_{i1} - \Phi_1}{\Phi_1} - (\ln \varphi_{i1} - \ln \Pi_1) \right) \\
&= \underbrace{\Delta s_i \left(\frac{\varphi_{i1} - \Pi_1}{\Pi_1} - (\ln \varphi_{i1} - \ln \Pi_1) \right)}_{\varepsilon_{i,B,appr}} + \underbrace{\Delta s_i \left(\frac{\varphi_{i1} - \Pi_1}{\Phi_1} - \frac{\varphi_{i1} - \Pi_1}{\Pi_1} \right)}_{\varepsilon_{i,B,ref}} + \underbrace{\Delta s_i \left(\frac{\Pi_1 - \Phi_1}{\Phi_1} \right)}_{\varepsilon_{i,B,\Delta mean}}
\end{aligned} \tag{12}$$

In eq. (12), we define the approximation error in the BFE ($\varepsilon_{i,B,appr}$) as the discrepancy between the log difference of $\ln \varphi_{i1}$ and $\ln \Pi_1$ and the ratio of $(\varphi_{i1} - \Pi_1)/\Pi_1$, which is approximated by this log difference. In line with Proposition (2), the reference productivity embodied in the log difference is Π_1 , while the reference productivity deployed by the BFE in levels is Φ_1 . This introduces the reference deviation ($\varepsilon_{i,B,ref}$). The remaining discrepancy is rooted in the mean deviation ($\varepsilon_{i,B,\Delta mean}$). In levels, firm productivity is measured against the arithmetic mean; for logs, it is compared to the geometric mean. The magnitude of $\varepsilon_{i,B,appr}$ increases with the difference $(\varphi_{i1} - \Pi_1)$ and is simultaneously scaled by the change of the firm's input share Δs_i . In line with Proposition (1), $\varepsilon_{i,B,appr}$ is positive for $\Delta s_i > 0$ and negative for $\Delta s_i < 0$. The impact of $\varepsilon_{i,B,ref}$ depends on the sign and magnitude of the change in input share Δs_i and the relative position of the individual firm's productivity φ_{i1} to the geometric mean Π_1 . Following Jensen's inequality, we assume that $\Phi_1 > \Pi_1$, so that $\varepsilon_{i,B,ref}$ is positive if $(\varphi_{i1} > \Pi_1 \wedge \Delta s_i < 0) \vee (\varphi_{i1} < \Pi_1 \wedge \Delta s_i > 0)$ and negative if $(\varphi_{i1} > \Pi_1 \wedge \Delta s_i > 0) \vee (\varphi_{i1} < \Pi_1 \wedge \Delta s_i < 0)$.

The size and magnitude of $\varepsilon_{i,B,\Delta mean}$ depend on the magnitude of the difference between the two means, $(\Pi_1 - \Phi_1)$, and the change in input share Δs_i . Provided that $\Phi_1 > \Pi_1$, the sign of the distortion will take the opposite sign of the change in input

share Δs_i , scaled by its absolute magnitude. In other words, for firms that increase their market share, the mean deviation induces a negative distortion in the BFE and vice versa. In the aggregate, the direction of the distortion caused by the mean deviation depends on the changes in market shares of surviving, exiting and entering firms. It will be positive if the sum of changes in shares is negative, i.e., $\sum_{i \in S} \Delta s_i < 0$, implying that the market share of surviving firms has decreased.

In sum, each of the three distortions in the BFE can be either positive or negative, both at the individual firm level and at the aggregate level. However, note that the aggregate impact of the mean deviation will occur only in an unbalanced panel. In a balanced panel, the aggregate impact of the mean deviation is always zero, as $(\Pi_1 - \Phi_1)/\Phi_1$ is constant and $\sum_{i \in S} \Delta s_i = 0$.

3.3. Log distortions in the cross-firm effect

The cross-firm effect (CFE) can be seen as an interaction between the two previous components, i.e., between the WFE and the BFE. From equations (9) and (10), we can derive the log distortions in the CFE:

$$\begin{aligned}
\varepsilon_{i,C} &= CFE_{i,lev} - CFE_{i,log} \\
&= \Delta s_i \left(\frac{\Delta \varphi_i}{\Phi_1} - \Delta \ln \varphi_i \right) \\
&= \underbrace{\Delta s_i \left(\frac{\Delta \varphi_i}{\varphi_{i1}} - \Delta \ln \varphi_i \right)}_{\varepsilon_{i,C,appr}} + \underbrace{\Delta s_i \left(\frac{\Delta \varphi_i}{\Phi_1} - \frac{\Delta \varphi_i}{\varphi_{i1}} \right)}_{\varepsilon_{i,C,ref}}
\end{aligned} \tag{13}$$

Again, we detect a log approximation error ($\varepsilon_{i,C,appr}$) that increases with $|\Delta \varphi_i|$; it is weighed and scaled by Δs_i and follows this scaling factor in sign and magnitude. The impact of the reference deviation ($\varepsilon_{i,C,ref}$) depends on the development of Δs_i , the relative position of φ_{i1} compared to Φ_1 , and $\Delta \varphi_i$. This variety of influencing factors leads to a large variety of potential outcomes concerning an over- or underestimation induced by logs.¹⁰

Overall, both the approximation error and the reference deviation may induce a positive or a negative bias in the CFE. It is also conceivable that the distortions balance out in the aggregate.

¹⁰For completeness: $\varepsilon_{i,C,ref}$ is positive for $(\varphi_{i1} > \Phi_1 \wedge \Delta s_i > 0 \wedge \Delta \varphi_i > 0)$ or $(\varphi_{i1} > \Phi_1 \wedge \Delta s_i < 0 \wedge \Delta \varphi_i < 0)$ or $(\varphi_{i1} < \Phi_1 \wedge \Delta s_i > 0 \wedge \Delta \varphi_i < 0)$ or $(\varphi_{i1} < \Phi_1 \wedge \Delta s_i < 0 \wedge \Delta \varphi_i > 0)$ and it is negative for all complementary cases.

3.4. Log distortions in entry and exit

The contribution of entering or exiting firms to aggregate productivity growth can be positive or negative. It depends on these firms' relative position to an industry's benchmark productivity, which, in the FHK decomposition, is the industry aggregate Φ_1 for levels and Π_1 for logs. Analogous to the BFE, we can again decompose the log-induced distortion to isolate three different distortions in both the entry and exit components: the log approximation error ($\varepsilon_{i,\cdot,appr}$), the reference deviation ($\varepsilon_{i,\cdot,ref}$), and the mean deviation ($\varepsilon_{i,\cdot,\Delta mean}$).

For entering firms, this leads to the following equation:

$$\begin{aligned}
 \varepsilon_{i,N} &= N_{i,lev} - N_{i,log} \\
 &= s_{i2} \left(\frac{\varphi_{i2} - \Phi_1}{\Phi_1} - (\ln \varphi_{i2} - \ln \Pi_1) \right) \\
 &= \underbrace{s_{i2} \left(\frac{\varphi_{i2} - \Pi_1}{\Pi_1} - (\ln \varphi_{i2} - \ln \Pi_1) \right)}_{\varepsilon_{i,N,appr}} + \underbrace{s_{i2} \left(\frac{\varphi_{i2} - \Pi_1}{\Phi_1} - \frac{\varphi_{i2} - \Pi_1}{\Pi_1} \right)}_{\varepsilon_{i,N,ref}} + \underbrace{s_{i2} \left(\frac{\Pi_1 - \Phi_1}{\Phi_1} \right)}_{\varepsilon_{i,N,\Delta mean}}
 \end{aligned} \tag{14}$$

The more distant the entering firm's productivity level φ_{i2} from the benchmark productivity level Π_1 is, the greater the log approximation error ($\varepsilon_{i,N,appr}$). In line with Proposition (1) and the fact that s_{i2} can take only positive values, $\varepsilon_{i,N,appr}$ is always non-negative. The distortion caused by the reference deviation ($\varepsilon_{i,N,ref}$) is almost identical to its counterpart in the BFE component, except for its weight. Assuming that $\Phi_1 > \Pi_1$, $\varepsilon_{i,N,ref}$ will be positive if $(\varphi_{i2} - \Pi_1) < 0$ and vice versa. The magnitude of the mean deviation ($\varepsilon_{i,N,\Delta mean}$) depends on the input share s_{i2} , which scales the difference between the geometric (Π_1) and arithmetic (Φ_1) means of productivity. Since the benchmark productivity for evaluating the contribution of entering firms is smaller for logs than for levels ($\Phi_1 > \Pi_1$), this distortion is always negative.

As is the case in the previous components, it is conceivable that the individual log distortions in the entry component balance out or induce an over- or underestimation, which depends on firm and industry characteristics.

In the case of exiting firms, we obtain the following mirror image:

$$\begin{aligned}
\varepsilon_{i,X} &= X_{i,lev} - X_{i,log} \\
&= s_{i1} \left(\frac{\Phi_1 - \varphi_{i1}}{\Phi_1} - (\ln \Pi_1 - \ln \varphi_{i1}) \right) \\
&= s_{i1} \left((\ln \varphi_{i1} - \ln \Pi_1) - \frac{\varphi_{i1} - \Phi_1}{\Phi_1} \right) \\
&= s_{i1} \underbrace{\left((\ln \varphi_{i1} - \ln \Pi_1) - \frac{\varphi_{i1} - \Pi_1}{\Pi_1} \right)}_{\varepsilon_{i,X,appr}} + s_{i1} \underbrace{\left(\frac{\varphi_{i1} - \Pi_1}{\Pi_1} - \frac{\varphi_{i1} - \Pi_1}{\Phi_1} \right)}_{\varepsilon_{i,X,ref}} + s_{i1} \underbrace{\left(\frac{\Phi_1 - \Pi_1}{\Phi_1} \right)}_{\varepsilon_{i,X,\Delta mean}}
\end{aligned} \tag{15}$$

The log approximation error in the Exit component ($\varepsilon_{i,X,appr}$) is always negative for exiting firms. The reference deviation ($\varepsilon_{i,X,ref}$) is positive for $\varphi_{i1} > \Pi_1$ and negative for $\varphi_{i1} < \Pi_1$. The mean deviation, ($\varepsilon_{i,X,\Delta mean}$) is always positive.

4. Empirical application

4.1. Data

We use firm-level panel data covering the French manufacturing sector for the 2009-18 period. The information comes from annual census data named *FARE* and covers more than 3 million companies per year. The data provide information about the firms' income statements and balance sheets, from which we retrieved data regarding value-added and the number of employees. Because we did not observe prices, we used industry-specific value-added deflators provided by the French statistical office INSEE. For labor, we used the industry-specific annual number of hours worked per employee (provided by INSEE) and multiplied it by the number of employees to obtain the total number of hours worked per company.

We restricted the analysis to manufacturing firms only.¹¹ Our motivation was to perform the decompositions on activities that were similar to what is generally documented in the literature. Moreover, we further restricted our sample to firms with at least 10 employees. Increasing the minimum size of the firm ensures higher data quality, a key element in growth rate computations. To avoid artificial breaks in the series, we did not trim observations with fewer than 10 employees on a firm-year basis. Rather, we screened out firms for which the median number of employees was strictly lower than 10 over the entire period.

We focused on labor productivity as our efficiency measure, defined as the value-

¹¹Table B1 of Appendix B lists the 12 two-digit manufacturing sectors. We excluded the industry of coke and refined petroleum products (ISIC 19) in our analysis.

added to hours worked ratio. We excluded firms reporting a negative value-added. We further truncated the data by excluding firms with at least one observation in the bottom and top .5% of the productivity distribution and by discarding firms that experienced suspicious negative and positive jumps in their efficiency series.¹² Applying such restrictions yielded a sample of approximately 260,000 firm-year observations. Table C1 of Appendix C reports the corresponding summary statistics.

In our empirical analysis, we followed two perspectives. First, we presented results for the average manufacturing industry by aggregating industry-level results to a manufacturing-wide level. Such aggregation exercises have been widely applied in productivity analyses with different degrees of detail (see e.g., Baily et al., 1992; Foster et al., 2001, 2006; Decker et al., 2017; Brown et al., 2018). Second, we investigated the productivity dynamics of individual industries, with a particular focus on the manufacturing of chemicals and chemical products (ISIC 20), whose productivity dynamics are well suited to illustrate the potential impact of log distortions.

Given our sample of firms, we performed the FHK decomposition using, alternatively, equations (9) and (10) for each industry and for each year. Because the decomposition exercise necessitates a starting year and an ending year, all results pertained to the period 2010-18, excluding 2009. This method yields a sample of $12 \times 9 = 108$ decompositions, allowing us to recover the overall log error ε_A ($\varepsilon_A = \sum_i \varepsilon_{i,A} = \hat{\Phi}_{lev} - \hat{\Phi}_{log}$) and decompose it into the error in the within component ε_W ($\varepsilon_W = \sum_i \varepsilon_{i,W}$), the error in the between component ε_B ($\varepsilon_B = \sum_i \varepsilon_{i,B}$), the error in the cross-firm effect component ε_C ($\varepsilon_C = \sum_i \varepsilon_{i,C}$), the error in the entry component ε_N ($\varepsilon_N = \sum_i \varepsilon_{i,N}$) and the error in the exit component ε_X ($\varepsilon_X = \sum_i \varepsilon_{i,X}$).

4.2. Log distortions in manufacturing

We reported the results for the manufacturing sector by creating a weighted average of the industry-level results for the 12 industries in our sample. As industry weights, we used labor input in the form of hours worked, averaged over the beginning and ending years of the period in which the respective growth rate was measured. The corresponding decomposition results and log distortions are shown in Table 1.

With regard to aggregate productivity growth, the table reports a distortion ε_A ranging from -0.64 to 1.36 percentage points. Relative to the productivity contribution in levels, the distortions ranged from -35% (2010) to 27% (2017). The log distortions in the aggregate growth rate were not distributed equally among the individual productivity components. The WFE and the exit and entry components were most affected, while the BFE and the CFE were subject to smaller distortions. Nonetheless, if expressed in relative terms, the distortions in the BFE and CFE can also be substantial. In all five components,

¹²More precisely, we excluded firms for which we observed a change in labor productivity by a factor of more than 3 from one year to another.

growth contributions were neither systematically overestimated nor underestimated by the use of logs. The distortions in the WFE ranged from -17% in 2010 to 25% in 2015. In the BFE component, the largest overestimation occurred in 2012 with -19% , and the largest underestimation occurred in 2011 with 18% . The range in the distortions in the CFE was from -14% in 2011 to 11% in 2018. The relative discrepancies were most pronounced with respect to the entry and exit components. For entries, the discrepancy ranged from -6% (2012) to 766% (2011). For exits, the distortions range from -239% in 2013 to 51% in 2018.

The evident lack of a systematic over- or underestimation in aggregate growth and in the components implies that trends in productivity developments may be judged differently in logs than in levels. For instance, calculated in levels, aggregate growth in 2017 was more than double that of 2010; in logs, it increased by only approximately 43% . As a further example, while aggregate growth in levels showed a slight increase between 2010 (1.83%) and 2014 (1.99%), logs suggested a slowdown in productivity growth from 2.47% to 1.83% . With respect to individual productivity components, the years 2016 and 2017 provided a very notable example of the potential misconceptions. Whereas with logs, the contributions of entering and exiting firms were very balanced, with levels, the exit component clearly dominated, being approximately three times the size of the entry component.

Moreover, we see that log distortions can lead to a flip of the sign of productivity components, i.e., turning a genuinely positive value into a negative value and vice versa. Over the reported period for the average industry, such a sign flip occurred in the entry component in 2011.

Apart from investigating the development of individual productivity components, decomposition methods are also frequently used to analyze which components have been the driving forces behind aggregate productivity growth over a given time span. As Table 1 sets out, such an analysis may be strongly blurred by logs. This is especially visible in the years 2010 and 2011, when evaluating the relevance of the BFE and the exit component. Both for levels and for logs, the WFE was the most relevant component for aggregate growth. However, while levels clearly point toward the BFE as the next most relevant component, for logs, the impact of exiting firms exceeded that of the BFE.

TABLE 1
Comparison of the decomposition results for levels and for logs

	WFE		BFE		CFE		Entry		Exit		Aggregate							
	lev	log	lev	log	lev	log	lev	log	lev	log	lev	log						
All manufacturing																		
2010	1.22	1.42	-0.21	0.83	0.82	0.01	-0.68	-0.65	-0.03	-0.09	-0.25	0.16	0.55	1.12	-0.57	1.83	2.47	-0.64
2011	2.29	2.27	0.02	0.92	0.76	0.16	-0.85	-0.74	-0.12	0.02	-0.15	0.18	0.65	1.01	-0.36	3.03	3.15	-0.12
2012	3.07	2.84	0.23	0.66	0.79	-0.13	-0.52	-0.47	-0.05	-0.67	-0.63	-0.04	-0.24	-0.19	-0.04	2.31	2.33	-0.03
2013	0.51	0.48	0.03	0.60	0.58	0.02	-0.41	-0.39	-0.03	-0.27	-0.65	0.39	0.12	0.41	-0.29	0.54	0.43	0.12
2014	1.70	1.49	0.21	0.46	0.46	0.00	-0.45	-0.44	-0.01	-0.23	-0.46	0.22	0.52	0.78	-0.27	1.99	1.83	0.16
2015	-1.44	-1.80	0.36	0.30	0.30	-0.01	-0.44	-0.42	-0.02	-0.07	-0.32	0.26	0.32	0.62	-0.30	-1.33	-1.61	0.28
2016	4.21	3.93	0.28	0.24	0.25	-0.01	-0.79	-0.69	-0.11	-0.13	-0.52	0.39	0.38	0.57	-0.18	3.91	3.53	0.37
2017	4.99	3.91	1.08	0.50	0.44	0.06	-0.63	-0.68	0.05	-0.07	-0.30	0.23	0.29	0.35	-0.06	5.07	3.72	1.36
2018	1.79	1.92	-0.14	0.42	0.48	-0.06	-0.62	-0.69	0.07	0.51	0.48	0.03	0.16	0.08	0.08	2.26	2.28	-0.02
Chemicals and chemical products																		
2010	-4.16	-3.53	-0.63	-0.09	-0.13	0.04	-0.36	-0.07	-0.29	1.41	1.03	0.38	-1.00	-0.26	-0.73	-4.20	-2.97	-1.23
2011	-1.76	-1.97	0.21	0.69	0.20	0.49	-0.89	-0.67	-0.22	2.29	1.81	0.48	0.38	0.41	-0.03	0.71	-0.22	0.94
2012	0.77	1.14	-0.36	0.47	0.37	0.11	-0.43	-0.27	-0.16	-0.29	-0.39	0.10	0.52	0.93	-0.41	1.05	1.77	-0.72
2013	7.94	8.62	-0.68	0.80	0.69	0.12	-0.84	-0.63	-0.21	-0.12	-0.29	0.17	-0.08	0.05	-0.13	7.71	8.44	-0.74
2014	-2.53	-3.00	0.47	0.35	0.46	-0.11	-0.80	-1.01	0.21	-0.14	-0.33	0.19	-0.37	-0.20	-0.16	-3.49	-4.08	0.59
2015	-5.10	-8.91	3.80	1.12	1.24	-0.12	-1.59	-1.73	0.14	-0.42	-0.65	0.23	0.25	0.34	-0.09	-5.75	-9.70	3.96
2016	15.05	16.89	-1.84	1.45	1.48	-0.03	-2.21	-2.27	0.06	-0.16	-0.24	0.08	-0.54	-0.20	-0.34	13.59	15.66	-2.07
2017	1.38	-1.26	2.63	-0.23	-0.01	-0.21	-0.12	-0.16	0.03	0.88	0.42	0.46	-0.02	0.20	-0.22	1.89	-0.81	2.70
2018	5.03	5.84	-0.81	0.51	0.48	0.03	-0.76	-0.71	-0.05	-0.09	-0.50	0.41	-0.49	-0.31	-0.18	4.20	4.80	-0.60

Notes: The top panel sets out the decomposition results and log distortions for 'All manufacturing', and the bottom panel shows results for the industry of 'Chemicals and chemical products' (ISIC 20). The results for 'All manufacturing' are based on the annual averages of the industry-level results for the 12 industries in our sample. As industry weights, we used labor input in the form of hours worked, averaged over the beginning and ending years of the period in which the respective growth rate was measured. The productivity components in levels and in logs are in %, the log distortions are in percentage points.

Looking at the general distribution of the log distortions, the summary statistics in Table 2 indicate an average positive tendency in the total distortions in the WFE (ε_W) and the entry component (ε_N), and a negative tendency in the distortion of the exit component (ε_X). The distortions for the BFE (ε_B) and CFE (ε_C) paint a rather balanced picture. Overall, this leads to a slightly positive tendency in the aggregate log distortion (ε_A), implying that, on average, logs underestimate aggregate productivity growth.

TABLE 2
Decomposition of the log distortions

		<i>All manufacturing</i>					<i>Chemicals and chemical products</i>				
		<i>Mean</i>	<i>Sd</i>	<i>Min</i>	<i>Med</i>	<i>Max</i>	<i>Mean</i>	<i>Sd</i>	<i>Min</i>	<i>Med</i>	<i>Max</i>
WFE	$\varepsilon_{W,appr}$	2.75	0.40	2.29	2.62	3.62	3.41	1.15	2.11	3.04	6.05
	$\varepsilon_{W,ref}$	-2.54	0.60	-3.83	-2.39	-1.79	-3.10	2.63	-7.89	-3.36	0.76
	ε_W	0.21	0.38	-0.21	0.21	1.08	0.31	1.80	-1.84	-0.36	3.80
BFE	$\varepsilon_{B,appr}$	0.14	0.18	-0.23	0.14	0.38	0.25	0.35	-0.13	0.21	1.08
	$\varepsilon_{B,ref}$	-0.06	0.03	-0.10	-0.06	-0.03	-0.08	0.07	-0.23	-0.08	0.01
	$\varepsilon_{B,\Delta Mean}$	-0.07	0.12	-0.26	-0.08	0.20	-0.14	0.16	-0.47	-0.10	0.09
	ε_B	0.00	0.08	-0.13	-0.00	0.16	0.04	0.20	-0.21	0.03	0.49
CFE	$\varepsilon_{C,appr}$	-0.03	0.08	-0.17	0.00	0.05	-0.08	0.46	-1.25	0.02	0.29
	$\varepsilon_{C,ref}$	0.00	0.09	-0.13	-0.01	0.22	0.03	0.51	-0.40	-0.12	1.31
	ε_C	-0.03	0.06	-0.12	-0.03	0.07	-0.05	0.18	-0.29	-0.05	0.21
Entry	$\varepsilon_{N,appr}$	0.54	0.12	0.36	0.55	0.75	0.61	0.33	0.25	0.57	1.16
	$\varepsilon_{N,ref}$	-0.02	0.04	-0.12	-0.01	0.01	-0.07	0.11	-0.28	-0.02	0.02
	$\varepsilon_{N,\Delta Mean}$	-0.31	0.20	-0.64	-0.20	-0.13	-0.26	0.16	-0.49	-0.21	-0.08
	ε_N	0.20	0.14	-0.04	0.22	0.39	0.28	0.16	0.08	0.23	0.48
Exit	$\varepsilon_{X,appr}$	-0.61	0.23	-1.06	-0.56	-0.35	-0.72	0.42	-1.50	-0.60	-0.18
	$\varepsilon_{X,ref}$	0.01	0.03	-0.03	-0.00	0.06	0.07	0.07	-0.01	0.05	0.19
	$\varepsilon_{X,\Delta Mean}$	0.38	0.16	0.23	0.39	0.74	0.39	0.25	0.10	0.36	0.86
	ε_X	-0.22	0.20	-0.57	-0.27	0.08	-0.26	0.21	-0.73	-0.18	-0.03
Aggr	$\varepsilon_{A,appr}$	2.78	0.37	2.36	2.63	3.34	3.47	0.66	2.69	3.44	4.61
	$\varepsilon_{A,ref}$	-2.61	0.64	-3.97	-2.45	-1.85	-3.16	2.36	-6.69	-3.57	0.52
	ε_A	0.16	0.53	-0.64	0.12	1.36	0.31	1.95	-2.07	-0.60	3.96

Notes: The table sets out the distribution of annual decomposed log distortions during the 2009-18 period. The left panel displays results for 'All manufacturing', and the right panel shows results for the industry of 'Chemicals and chemical products' (ISIC 20). The results for 'All manufacturing' are based on the annual averages of industry-level results for the 12 industries in our sample. As industry weights, we used labor input measured by hours worked, averaged over the beginning and ending years of the period in which the respective growth rate was measured. All log distortions are reported in percentage points.

The decomposition of the log distortions in Table 2 offers an explanation for the observed tendencies. With respect to the WFE component, the approximation error ($\varepsilon_{W,appr}$) and the reference deviation ($\varepsilon_{W,ref}$) appeared to offset each other to a certain extent. The approximation error was positive with a mean of 2.75 percentage points and thereby exceeded the reference deviation, which was consistently negative with a mean of -2.54 percentage points. When comparing the range and the standard deviation of the two distortions, it appeared to be the reference deviation causing the fluctuations in the aggregate distortion of the WFE (ε_W). The negative mean of the reference deviation is an artifact, as we cannot conclude the sign of the aggregate distortion of $\varepsilon_{W,ref}$ from the

calculations in Section 3. Only 3 out of the 108 industry-year combinations the average industry is based on showed a positive reference deviation. A possible explanation of the negative tendency is that firms with below-average productivity ($\varphi_{i1} < \Phi_1$) tend to increase their productivity, whereas firms with above-average productivity ($\varphi_{i1} > \Phi_1$) tend to decrease their productivity.¹³ Overall, this contributes to a mostly negative reference deviation in the WFE (compare Section 3.1).

Regarding the BFE component, the average total log distortion (ε_B) was approximately zero. Decomposing this distortion, the reference deviation ($\varepsilon_{B,ref}$) was consistently negative and showed little volatility compared to the other two error terms. The approximation error ($\varepsilon_{B,appr}$) and mean deviation ($\varepsilon_{B,\Delta mean}$) showed years of positive and negative distortions. On average, the positive tendency in the approximation error and the negative tendency in the mean deviation can be attributed to the individual firm/industry characteristics of our sample and may differ for another sample. The negative median in the mean deviation, for instance, shows that the years in which surviving firms increased their market share compared to entering and exiting firms outnumber the years with a decrease in surviving firms' market share (compare Section 3.2).

The decomposition of the log distortion in the CFE (ε_C) reveals a slightly negative tendency in the approximation error ($\varepsilon_{C,appr}$), while the reference deviation ($\varepsilon_{C,ref}$) averages approximately zero. As highlighted in the previous section, the direction of the two distortions depends on the direction of both the change in market share and the change in productivity. Due to this variety in influencing factors, neither our theoretical nor our empirical approach revealed any discernible pattern within the two distortions.

The positive tendency in the log distortion in the entry component (ε_N) was mainly caused by the approximation error ($\varepsilon_{N,appr}$), which was consistently positive (compare Section 3.4). Even though the mean deviation ($\varepsilon_{N,\Delta mean}$) was consistently negative, Table 2 shows that the approximation error was dominant. The reference deviation ($\varepsilon_{N,ref}$) seemed to be negligible in most years.

The reasoning for the negative tendency of the log distortion in the exit component (ε_X) mirrors that of the entry component. The approximation error ($\varepsilon_{X,appr}$) was consistently negative, and the mean deviation ($\varepsilon_{X,\Delta mean}$) was consistently positive but, on average, was surpassed by the approximation error. The reference deviation ($\varepsilon_{X,ref}$) played only a minor role in most years.

¹³Based on our industry-level results for 2009-18, 57% of firms with below-average productivity increased productivity, and 58% of above-average firms experienced a decrease in productivity (averages over all industries and years). The shares alone, however, are only an indication because the reference deviation is also affected by the difference between firm-level productivity and the mean reference, the magnitude of productivity growth rates, and the input shares of firms in the respective category (see Section 3).

4.3. Log distortions in individual industries

We now turn to log distortions when investigating firm-level productivity dynamics within individual industries. As mentioned above, we applied equations (9) and (10) to each of the 12 industries listed in Table B1 and subsequently calculated the log-induced discrepancy, resulting in 108 industry-year observations. We treated the large quantity of results by reporting percentiles in Table 3. Subsequently, we present detailed results for manufacturers of chemicals and chemical products (ISIC 20). Over the 2009-18 period, this industry was shaped by strong productivity growth, which highlights the potential misconceptions induced by logs in productivity decompositions.

TABLE 3
Log distortions in individual industries. N = 108

	<i>Min</i>	p_{10}	p_{25}	<i>Med</i>	p_{75}	p_{90}	<i>Max</i>
ε_W	-1.84	-0.78	-0.29	0.24	0.66	1.46	3.80
ε_B	-0.56	-0.17	-0.07	-0.01	0.08	0.15	0.77
ε_C	-0.85	-0.23	-0.09	-0.03	0.03	0.10	0.49
ε_N	-1.30	0.02	0.08	0.16	0.32	0.52	1.22
ε_X	-2.94	-0.73	-0.36	-0.20	-0.10	-0.04	1.44
ε_A	-3.36	-1.15	-0.43	0.09	0.72	1.34	4.56

Notes: The table sets out the distribution of the annual log distortions occurring in the decomposition exercises conducted for the 12 industries in our sample in the period 2009-18. For each component, the number of industry-year combinations is $N = 108$. The reported values for the log distortions are in percentage points; $p_{(\cdot)}$ reflect percentiles.

As shown in Table 3, most distortions stayed within the range of approximately ± 1 percentage points, although it is possible that log distortions reached or even exceeded values of 3 percentage points. Analogous to our previous findings, the BFE and the CFE were, in absolute terms, less affected by log distortions. Nonetheless, even these two components could be subject to considerable distortions, exceeding our observations for the average industry. Moreover, for the 108 industry-year combinations within our sample, we identified 25 combinations with a sign flip either in at least one of the five components or in aggregate productivity growth. This result implies that almost one out of four industry-year combinations was affected by a sign flip. This result underscores the potential impact of log distortions on individual industries.

Comparing the distortions between different industries, we detected that some industries were strongly affected by logs, while others were less affected. This means that when performing a decomposition exercise for an individual industry, the results are not necessarily strongly distorted by logs. While we argued in Section 3 that log distortions essentially depend on the features and development of individual firms within a given industry, the observed differences here raise the question of whether there is a systematic pattern or certain characteristics that make an industry more or less prone to log distortions.

Before we investigate this relationship in Section 5, we turn to the decomposition re-

sults for manufacturers of chemicals and chemical products (Table 1). Overall, the results are in line with our findings regarding the average industry. However, the magnitude and fluctuations of the distortions exceeded those for the average industry. The distortions in aggregate productivity growth ranged from -2.07 to 3.96 percentage points. In relative terms, the distortion varied between -69% (2012) and 143% (2017). These distortions were mostly driven by the large deviations in the WFE, which ranged from -47% in 2012 to 191% in 2017. In the BFE component, the span of distortions reached from -94% in 2017 to 71% in 2011. For the CFE, the largest overestimation amounted to -81% in 2010, while the underestimation was the strongest in 2017, with 28% . Once again, the relative distortions were most pronounced in the entry and exit components. In the case of the entry component, the discrepancy was always positive, ranging from 21% in 2011 to 434% in 2018. The distortion in the exit component was consistently negative, ranging from -1122% in 2017 to -8% in 2011.

It is obvious that the combination of the magnitude and volatility of these log distortions could lead to severe misconceptions concerning productivity growth. This becomes especially evident in the prevalence of sign flips in aggregate productivity growth (2011 and 2017), the WFE (2017) and the exit component (2013 and 2017), as reported in Table 1. A look at the BFE component reveals a further example of a misconception. Considering the development of the BFE between 2011 and 2012, logs created the impression that the BFE had almost doubled between 2011 and 2012, whereas the results in levels showed that the BFE had actually decreased.

The described high fluctuations are also reflected in the large standard deviation of the log distortions in all components, especially in the WFE, BFE, CFE, and aggregate productivity growth (Table 2). Despite the volatility in distortions, we again detected an average positive distortion in the WFE (ε_W), the entry component (ε_N), and aggregate growth (ε_A), whereas the average distortion in the exit component (ε_X) was negative. What was most striking in the results for the chemicals industry in Table 2 was the occurrence of positive values in the reference deviation of the WFE ($\varepsilon_{W,ref}$). Recall that the reference deviation exhibited a strong negative tendency in our sample, for which we offered the opposite direction of the development of below- and above-average productivity firms as a possible explanation. For the chemicals industry, we detected slightly positive values for the reference deviation in the years 2015 (0.76) and 2017 (0.064). Again, the development of firms with below- and above-average productivity may provide one possible explanation for the reference deviation in these two years. Since both below- and above-average firms mostly decreased their productivity in 2015 and 2017, the negative reference deviation of the first group seems to be compensated for by the positive reference deviation of the second group, yielding an overall positive reference deviation.¹⁴

¹⁴Of firms with an initial productivity above the mean ($\varphi_{i1} > \Phi_1$), 64% (2015) and 61% (2017) decreased their productivity. Of below-average firms ($\varphi_{i1} < \Phi_1$), 61% and 55%, respectively, decreased their productivity.

Hence, by revealing how log distortions are driven by the idiosyncratic development of firms and industries, the chemicals industry offers an instructive example of the difficulty in predicting the impact of log distortions.

In summary, our results show that logs do induce distortions but not in a *systematic* manner. Consequently, the conclusions drawn must lead to fallacies regarding the growth contribution of the individual productivity components as well as in overall productivity growth. In our empirical exercise, we found that the use of logs led to an underestimation of the WFE and the entry component, while the exit component was overestimated. Overall, this caused an underestimation of total productivity growth. The BFE and the CFE were less affected; in relative terms, however, they could still be subject to large distortions. With respect to the decomposed log distortions themselves, we found that all three identified log distortions (approximation error, reference deviation, mean deviation) contributed decisively to the discrepancy between level and log components, albeit with different variations. Despite the numerous parallels between the average and individual industries, there were differences in the degree to which logs affected the decomposition results. Although the use of logs may affect some industries more than others, the magnitude of potential distortions on the individual industry level was significantly larger than that for the aggregate average, a phenomenon that is well known as the fallacy of compound (composition): on average, overestimates and underestimates tend to cancel each other out and thus (partially) mask log-induced inaccuracies at the industry level.

5. The fallacy in decomposition and industry characteristics

The previous section showed that log-based decompositions embody what we call the fallacy in decomposition: using logs may lead to either an inaccurate aggregate growth rate, an inaccurate description of the contribution of the microsources, or both. In this section, we quantify this fallacy in order to have a general idea of the accuracy of log-based decompositions. Moreover, we show that log distortions correlate in a systematic fashion with certain industry characteristics.

To simplify our investigation, we defined three microsources of economic growth: firm learning, defined as the within firm effect holding the input share constant and allowing firm efficiency to vary (WFE in equations 9 and 10); the resource reallocation effect, resulting from changes in the input shares of firms ($BCFE = BFE + CFE$ in equations 9 and 10); and industry churning, defined as the effect of entry into and exit from the market ($NXE = N - X$). Table 4 presents the four possible hypothetical cases that may arise when using logs, as in eq. (10). It is based on the assumption that the correct decomposition exercise is performed using levels according to eq. (9). In case 1, the log-based decomposition produces the same results as those obtained from the level-based

decomposition. In case 2, using logs leads to an overestimation of aggregate growth, while leaving the contribution (expressed in %) of each component unaffected. In case 3, log-based decomposition generates inaccurate contributions while leaving aggregate growth unaffected. Case 4 is the worst-case scenario when both types of mismeasurement apply. Cases 2 to 4 all represent what we call the fallacy in decomposition when using logs. All imply an incorrect representation of the sources of aggregate productivity growth.

TABLE 4
Decomposition fallacy: four hypothetical cases

	<i>AG</i>	<i>WFE</i>	<i>BCFE</i>	<i>NXE</i>
<i>Results using level-based decomposition</i>				
True values	4 <i>100</i>	2 <i>50</i>	1.5 <i>37.5</i>	0.5 <i>12.5</i>
<i>Possible cases using log-based decomposition</i>				
Case 1: No error	4 <i>100</i>	2 <i>50</i>	1.5 <i>37.5</i>	0.5 <i>12.5</i>
Case 2: Error in aggregate growth	6 <i>100</i>	3 <i>50</i>	2.25 <i>37.5</i>	0.75 <i>12.5</i>
Case 3: Error in contributions	4 <i>100</i>	1 <i>25</i>	2 <i>50</i>	1 <i>25</i>
Case 4: Both types of errors	6 <i>100</i>	1.5 <i>25</i>	3 <i>50</i>	1.5 <i>25</i>

Notes: Figures in italics represent the respective contributions of the within-firm effect (*WFE*), the reallocation effect (*BCFE*), and the net-entry effect (*NXE*) in % of aggregate growth (*AG*). In line with the definition we gave in the introduction, we consider the level-based decomposition to provide the 'true values'.

To quantify the fallacy in decomposition, we analyzed whether logs led to an inaccurate description of aggregate growth and of the microfoundations of productivity growth. To do so, we performed the FHK decomposition using, alternatively, equations (9) and (10) at the four-digit industry level for each year. We then measured the overall log distortion ε_A ($\varepsilon_A = \sum_i \varepsilon_{i,A}$) and the log distortions in the individual productivity components, i.e., in the within component ε_W ($\varepsilon_W = \sum_i \varepsilon_{i,W}$), the reallocation component ε_{BC} ($\varepsilon_{BC} = \sum_i \varepsilon_{i,B} + \sum_i \varepsilon_{i,C}$), and the churning component ε_{NX} ($\varepsilon_{NX} = \sum_i \varepsilon_{i,N} + \sum_i \varepsilon_{i,X}$). Trimming the bottom and the top 1% of each of these log distortions ($\varepsilon_A, \varepsilon_W, \varepsilon_{BC}, \varepsilon_{NX}$) yielded 2,805 observations, of which we could retrieve all information.

We infer the fallacy in decomposition by simply counting the frequency of 'inaccurate' measurements, which we define as the following. First, with respect to aggregate growth, we arbitrarily qualify a measurement as 'accurate' if aggregate growth in logs does not differ from aggregate growth in levels by more than $\pm\alpha\%$, where α represents the tolerance level below which the log measure is considered accurate: $|\varepsilon_A/\hat{\Phi}_{lev}| \leq \alpha\%$. Second, we define the measurement of the contribution of components as accurate if no log components' contribution to aggregate growth in logs deviates by more than α percentage points from the respective counterpart in levels. That is: $\Delta\theta_W \wedge \Delta\theta_{BC} \wedge \Delta\theta_{NX} < \alpha/100$, where

$\Delta\theta_Z = Z_{lev}/\hat{\Phi}_{lev} - Z_{log}/\hat{\Phi}_{log}$ and $Z = \{W; BC; NX\}$. We then counted the frequency of accurate and inaccurate measurements of both aggregate growth and the contribution of the individual productivity components and built a 2×2 table presenting the frequencies of the four possible cases displayed in Table 4.

TABLE 5
The fallacy in decomposition – four-digit industry level. $N = 2,805$

<i>Aggregate growth</i>	<i>Contribution of components</i>	
	Accurate	Inaccurate
$\alpha = 5$		
Accurate	7.9	6.5
Inaccurate	23.7	61.9
$\alpha = 10$		
Accurate	21.6	6.6
Inaccurate	28.2	43.6
$\alpha = 20$		
Accurate	43.3	5.9
Inaccurate	24.8	26.0

Notes: Parameter α denotes the tolerance level that determines whether a measure is accurate or inaccurate. As for aggregate growth, it represents the magnitude of the aggregate distortion (ε_A) relative to the true growth rate ($\hat{\Phi}_{lev}$). A log-based aggregate growth rate is considered accurate when $|\varepsilon_A/\hat{\Phi}_{lev}| \leq \alpha\%$. The contribution of components is considered accurate if no log components' contribution to aggregate growth in logs deviates by more than α percentage points from the respective level counterpart. That is: $\Delta\theta_W \wedge \Delta\theta_{BC} \wedge \Delta\theta_{NX} < \alpha/100$, where $\Delta\theta_Z = Z_{lev}/\hat{\Phi}_{lev} - Z_{log}/\hat{\Phi}_{log}$ and $Z = \{W; BC; NX\}$. The χ^2 test reveals that the two events 'accuracy in aggregate growth using logs' and 'accuracy of contributions using logs' are related at 1% significance level, irrespective of the tolerance level α . Numbers in the table are in % of industry-year combinations, that is, of $N = 2,805$.

Table 5 presents the results for three tolerance levels: $\alpha = 5$, $\alpha = 10$, and $\alpha = 20$. Starting with a low tolerance level where $\alpha = 5$, we observed that for only 8% of the decomposition exercises, the log-based decomposition exercise proved accurate. In the majority of cases (62%), decomposition using logs yielded an inaccurate aggregate growth rate and inaccurate contributions of the three components. In almost one case in four, log-based decomposition yielded inaccurate aggregate growth rates without affecting individual contributions. This result implies that, in these cases, the overall mismeasurement in aggregate growth stemmed from a roughly equal mismeasurement in all growth components. Only 6.5% of the decomposition exercises yielded inaccurate contributions with accurate aggregate growth rates. Altogether, at $\alpha = 5$, decompositions were inaccurate in more than 9 out of 10 cases.

Increasing the tolerance level α to 10 and 20 mechanically increased the number of accurate decomposition exercises to 22% and 43%, respectively. However, this barely affected the off-diagonal elements, amounting to approximately 3 cases in 10. This result implies that increasing the tolerance level barely affects the fact that 30% of such log-based decompositions produce inaccurate results. The second observation was that although an increase in the tolerance level increased the frequency of accurate results, a substantial

number of decompositions remained inaccurate: 8 decompositions out of 10 for $\alpha = 10$ and 6 out of 10 for $\alpha = 20$. Altogether, the message from Table 5 is that log-based decompositions generally yield inaccurate results, and accurate decompositions mostly represent the exception, not the rule.

The consequences of these inaccuracies also affect the validity of studies that investigate the role of industry characteristics in productivity growth. To draw a rough sketch of the impact of log-induced decomposition fallacies in such ventures, we investigated the extent to which log-induced distortions were associated with industry characteristics. We deployed industry characteristics often used in the literature as a candidate explanation for the observed aggregate productivity growth. These comprehend export intensity (ExpInt: industry sum of export divided by the industry sum of sales), profit rate (PrRate: industry sum of profit divided by the industry sum of value-added), investment rate (InvRate: industry sum of investment divided by the industry sum of value-added), the number of firms in the four-digit industry (firm count FC, in logs), mean firm size (MFS, in logs: industry sum of working hours divided by industry number of firms), and industry concentration as measured by the Herfindahl-Hirschman Index for sales as market shares (HHI). We did not have any particular prior on whether and how these industry characteristics were associated with log distortions, and by no means do we intend to depict causal relationships running from industry characteristics to log distortions. This advocates the use of an ordinary least squares estimator and the model specification reads

$$\mathbf{Y}_{st} = \alpha + \mathbf{B}'\mathbf{X}_{st} + \epsilon_{st} \quad (16)$$

where $\mathbf{Y} = \{\varepsilon_A, \varepsilon_W, \varepsilon_{BC}, \varepsilon_{NX}\}$ and \mathbf{X} includes the six industry characteristics mentioned above. Subscripts s and t stand for four-digit sector s at time t . Column vector \mathbf{B} represents the parameter estimates, which in this case, should be interpreted as mere partial correlation coefficients.

Documenting how industry characteristics are associated with log distortions is not as straightforward as it may initially seem. Since log distortions can be positive or negative, the signs of the parameter estimates in model (16) cannot simply be interpreted as increasing or decreasing the distortion. To illustrate, imagine that the distortion is positive ($\mathbf{Y} > 0$), i.e., that logs underestimate the productivity component. Then, a positive parameter estimate suggests that the given industry characteristic is positively associated with log distortions. Instead, imagine that the average distortion is negative. Then, a positive parameter implies that the given industry characteristic moderates log distortions. To resolve this ambiguity, one could choose to use the absolute value of distortions as the LHS variable. This, however, excludes the possibility of an asymmetrical correlation, i.e., that a given industry characteristic only underestimates but does not overestimate aggregate growth and vice versa. Therefore, we allowed the partial correlations to differ between a positive ($\mathbf{Y} > 0$) and a negative log distortion ($\mathbf{Y} < 0$) and ran model (16) on

the two respective subsamples.

Furthermore, we built our two subsamples exclusively on whether ε_A was positive or negative. In our data, the overall number of observations was 2,805, of which 1,619 pertained to an underestimation of aggregate productivity growth ($\varepsilon_A > 0$), and 1,186 pertained to an overestimation of aggregate productivity growth ($\varepsilon_A < 0$). We then regressed ε_A , ε_W , ε_{BC} and ε_{NX} sequentially on the vector of explanatory variables \mathbf{X} . Because $\varepsilon_A = \varepsilon_W + \varepsilon_{BC} + \varepsilon_{NX}$, the reported parameter estimates pertaining to the dependent variables ε_W , ε_{BC} and ε_{NX} all sum to the estimate pertaining to ε_A : $\hat{\beta}_{\varepsilon_A} = \hat{\beta}_{\varepsilon_W} + \hat{\beta}_{\varepsilon_{BC}} + \hat{\beta}_{\varepsilon_{NX}}$. This method allows us to depict where the sources of the overall log distortion stem from and whether this affects the respective contributions.¹⁵

The left (right) panel of Table 6 displays the results for underestimated (overestimated) aggregate growth rates. Focusing first on the left panel and starting with export intensity, we observed that sectors more committed to international trade were associated with larger log-induced underestimations ($\hat{\beta}_{\varepsilon_A, ExpInt}^{\varepsilon_A > 0} = 0.615$). This result mainly stems from a significant underestimation of the within component ($\hat{\beta}_{\varepsilon_W, ExpInt}^{\varepsilon_A > 0} = 0.872$), though partially compensated by a moderating churning coefficient ($\hat{\beta}_{\varepsilon_{NX}, ExpInt}^{\varepsilon_A > 0} = -0.310$). This result implies that the contribution of firm learning is systematically underestimated in more open industries. Looking at the right panel, we found no significant overestimation

¹⁵This also implies that the interpretation of whether the respective industry characteristic exacerbates or reduces log distortions applies to ε_A exclusively.

TABLE 6
Industry characteristics and the magnitude of log distortions using logs

	<i>Log-induced underestimation ($\varepsilon_A > 0$)</i>				<i>Log-induced overestimation ($\varepsilon_A < 0$)</i>			
	ε_A	ε_W	ε_{BC}	ε_{NX}	ε_A	ε_W	ε_{BC}	ε_{NX}
ExpInt	0.615** (0.262)	0.872*** (0.249)	0.053 (0.078)	-0.310* (0.175)	0.114 (0.254)	0.395* (0.228)	-0.091 (0.082)	-0.190 (0.211)
PrRate	0.912** (0.379)	0.279 (0.361)	0.500*** (0.113)	0.134 (0.253)	-0.041 (0.366)	0.651** (0.328)	0.033 (0.118)	-0.726** (0.305)
HHI	3.306*** (0.608)	1.828*** (0.580)	0.308* (0.181)	1.171*** (0.406)	1.785*** (0.554)	1.786*** (0.497)	0.039 (0.179)	-0.041 (0.461)
InvRate	0.005 (0.060)	-0.003 (0.058)	0.019 (0.018)	-0.011 (0.040)	-0.110* (0.062)	-0.095* (0.056)	0.003 (0.020)	-0.018 (0.052)
FC	-0.271*** (0.047)	-0.250*** (0.044)	0.026* (0.014)	-0.046 (0.031)	-0.307*** (0.043)	-0.143*** (0.039)	-0.007 (0.014)	-0.158*** (0.036)
MFS	-0.397*** (0.073)	-0.257*** (0.070)	-0.014 (0.022)	-0.126** (0.049)	-0.193*** (0.066)	-0.036 (0.059)	-0.018 (0.021)	-0.139** (0.055)
Observations	1,619	1,619	1,619	1,619	1,186	1,186	1,186	1,186
R-squared	0.110	0.083	0.017	0.014	0.117	0.093	0.002	0.028

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Constant not reported. The left panel shows log-induced underestimations ($\varepsilon_A > 0$), and the right panel depicts overestimations ($\varepsilon_A < 0$). The dependent variable is the absolute value of the difference between levels and logs for the individual components: Aggregate (ε_A); Within (ε_W); Reallocation (ε_{BC}); and Churning (ε_{NX}). A positive (resp. negative) sign for the parameter estimate indicates that a given industry characteristics is associated with a larger (resp. smaller) under- or overestimation using logs.

issue for more open industries, except for the within component ($\hat{\beta}_{\varepsilon_W, ExpInt}^{\varepsilon_A < 0} = 0.395$). Altogether, industry openness is associated with a systematic underestimation of aggregate growth and affects the contributions of the microsources of growth.

Turning to the profit rate, we observed a similar pattern: industries with higher profit rates were associated with a larger log-induced underestimation ($\hat{\beta}_{\varepsilon_A, PrRate}^{\varepsilon_A > 0} = 0.912$), while there was no significant association with an overall log-induced overestimation. In the former case, the distortions stemmed essentially from an underestimation of the reallocation component ($\hat{\beta}_{\varepsilon_{BC}, PrRate}^{\varepsilon_A > 0} = 0.500$); in the latter case, where no significant association between profit rate and a log-induced overestimation could be identified, the distortion was caused by a positive association of the within component ($\hat{\beta}_{\varepsilon_{BC}, PrRate}^{\varepsilon_A < 0} = 0.651$), which was compensated by the negative association of the churning coefficient ($\hat{\beta}_{\varepsilon_{NX}, PrRate}^{\varepsilon_A < 0} = -0.726$).

Industry concentration is a characteristic that always exacerbates log distortions by increasing both under- and overestimations ($\hat{\beta}_{\varepsilon_A, HHI}^{\varepsilon_A > 0} = 3.306$ and $\hat{\beta}_{\varepsilon_A, HHI}^{\varepsilon_A < 0} = 1.785$). Such distortions spread across all components, except for the reallocation and the net entry distortion component in the case of overestimation. This is a clear indication that highly concentrated industries are prone to distortions when decomposing aggregate productivity growth based on log-transformed measures of efficiency. This not only affects the estimated productivity growth but also casts doubt on the relevance of the contribution of each component.

Looking at the investment rate, industries with a high investment rate appeared not to correlate with log distortions (apart from a negligible but still significant distortion-reducing firm learning effect in the case of overestimation), neither concerning the aggregate growth rate nor by modifying the respective contributions of the microsources of growth.

The two key variables that significantly reduced log distortions, by reducing both under- and overestimations, were the number of firms and the average size of firms. Regarding the firm count, this goes through a reduction in the distortions in the within and churning components, ε_W and ε_{NX} . For the reallocation component, in contrast, a high number of firms was associated with a larger underestimation ($\hat{\beta}_{\varepsilon_{BC}, FC}^{\varepsilon_A > 0} = 0.026$). In fact, increasing the number of firms in an industry is tantamount to increasing the level of competition. Increased competition should lead to a reduction in profit rates, lower market shares and lower price-cost margins. This, in turn, should translate into less right-skewed distributions in sales and size, with the presence of dominant firms being undermined. Concerning average firm size, the decrease in distortions stems from a reduction in the churning component across both over- and underestimations, supported by a distortion-reducing firm-learning component. Our interpretation is that a higher mean firm size is a proxy for entry barriers. In turn, fewer movements in firm entry and exit reduce distortions due to industry churning.

Altogether, the results unambiguously show that variables proxying for lack of competition, such as industry concentration or profit rate, and for the openness of the industry, such as export intensity, are associated with log distortions. Conversely, more competitive industries are associated with lower log distortions.

6. Conclusion

The use of logs in productivity decomposition induces fallacious conclusions: using logs may lead to either inaccurate aggregate productivity growth, an inaccurate description of the contribution of the productivity components, or both. As we show in our paper, this fallacy is due to three log distortions: (i) the log approximation error, as a consequence of the logarithm's concavity; (ii) the reference deviation, arising from a different reference assumption implicit in log differences; and (iii), the mean deviation, caused by the difference in the deployed benchmark productivity.

Carried out on the basis of the FHK decomposition method, we calculated the respective distortions analytically and showed their magnitude empirically using firm-level data of the French manufacturing sector during 2009-18. The results suggest that the use of logs can lead to substantial misconceptions regarding productivity developments. Log-induced distortions appear to be unsystematic, which implies that each productivity component as well as aggregate productivity growth may be either over- or underestimated. This impairs the comparison between log and level results as well as the comparison between log results themselves. Overall, however, our empirical exercise suggests that logs underestimate the growth contribution stemming from both the within-firm and the entry component, while overestimating the contribution of the exit component; conversely, the between- and cross-firm components appear to be less affected by the use of logs. In sum, these tendencies caused and underestimation of aggregate productivity growth in our study.

Performing decompositions at a fine-grained industry level has allowed us to quantify this fallacy in log-based decompositions. As the results show, even with reasonably high levels of tolerance, the odds are high that a log-based decomposition will yield misleading results. With a simple study on the association of industry characteristics with log distortions, we further show that the magnitude of log distortions is substantial for inferential productivity analyses: on the aggregate level as well as on the level of productivity components, log distortions correlate significantly with industry characteristics. Consequently, the use of logs will inevitably induce severe endogeneity problems in inferential regression analyses when using log-transformed productivity components.

References

- Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., and Kerr, W. (2018). Innovation, reallocation, and growth. *American Economic Review*, 108(11):3450–91.
- Alon, T., Berger, D., Dent, R., and Pugsley, B. (2018). Older and slower: The startup deficit’s lasting effects on aggregate productivity growth. *Journal of Monetary Economics*, 93:68–85.
- Baily, M. N., Bartelsman, E. J., and Haltiwanger, J. (1996). Downsizing and productivity growth: myth or reality? *Small Business Economics*, 8(4):259–278.
- Baily, M. N., Bartelsman, E. J., and Haltiwanger, J. (2001). Labor productivity: structural change and cyclical dynamics. *Review of Economics and Statistics*, 83(3):420–433.
- Baily, M. N., Hulten, C., and Campbell, D. (1992). Productivity dynamics in manufacturing plants. *Brookings Papers on Economic Activity. Microeconomics*, 1992:187–267.
- Bartelsman, E., Haltiwanger, J., and Scarpetta, S. (2004). Microeconomic evidence of creative destruction in industrial and developing countries. *Available at SSRN 612230*.
- Bartelsman, E., Haltiwanger, J., and Scarpetta, S. (2009). Measuring and analyzing cross-country differences in firm dynamics. In *Producer Dynamics: New Evidence from Micro Data*, pages 15–76. University of Chicago Press.
- Bartelsman, E., Haltiwanger, J., and Scarpetta, S. (2013). Cross-country differences in productivity: The role of allocation and selection. *American Economic Review*, 103(1):305–34.
- Bernard, A. B., Eaton, J., Jensen, J. B., and Kortum, S. (2003). Plants and productivity in international trade. *American Economic Review*, 93(4):1268–1290.
- Brown, J. D., Crespi, G. A., Iacovone, L., and Marcolin, L. (2018). Decomposing firm-level productivity growth and assessing its determinants: evidence from the americas. *The Journal of Technology Transfer*, 43(6):1571–1606.
- Casella, G. and Berger, R. L. (2002). Statistical inference. duxbury. *Pacific Grove, CA*.
- Davis, S. J. and Haltiwanger, J. (1999). Gross job flows. *Handbook of Labor Economics*, 3:2711–2805.
- Decker, R. A., Haltiwanger, J., Jarmin, R. S., and Miranda, J. (2017). Declining dynamism, allocative efficiency, and the productivity slowdown. *American Economic Review*, 107(5):322–26.
- Dias, D. A. and Marques, C. R. (2021a). From micro to macro: a note on the analysis of aggregate productivity dynamics using firm-level data. *Journal of Productivity Analysis*, pages 1–14.
- Dias, D. A. and Marques, R. C. (2021b). Every cloud has a silver lining: Cleansing effects of the portuguese financial crisis. *Oxford Bulletin of Economics and Statistics*, 83(2):352–376.
- Disney, R., Haskel, J., and Heden, Y. (2003). Restructuring and productivity growth in uk manufacturing. *The Economic Journal*, 113(489):666–694.

- Dosi, G. and Grazzi, M. (2006). Technologies as problem-solving procedures and technologies as input–output relations: some perspectives on the theory of production. *Industrial and Corporate Change*, 15(1):173–202.
- Dosi, G., Moschella, D., Pugliese, E., and Tamagni, F. (2015). Productivity, market selection, and corporate growth: comparative evidence across us and europe. *Small Business Economics*, 45(3):643–672.
- Fagerberg, J. (2000). Technological progress, structural change and productivity growth: a comparative study. *Structural Change and Economic Dynamics*, 11(4):393–411.
- Färe, R. and Karagiannis, G. (2017). The denominator rule for share-weighting aggregation. *European Journal of Operational Research*, 260(3):1175–1180.
- Foster, L., Haltiwanger, J., and Krizan, C. J. (2006). Market selection, reallocation, and restructuring in the us retail trade sector in the 1990s. *The Review of Economics and Statistics*, 88(4):748–758.
- Foster, L., Haltiwanger, J. C., and Krizan, C. J. (2001). Aggregate productivity growth: lessons from microeconomic evidence. In *New Developments in Productivity Analysis*, pages 303–372. University of Chicago Press.
- Griliches, Z. and Regev, H. (1995). Firm productivity in israeli industry 1979–1988. *Journal of Econometrics*, 65(1):175–203.
- Hakkala, K. (2006). Corporate restructuring and labor productivity growth. *Industrial and Corporate Change*, 15(4):683–714.
- Haskel, J. and Sadun, R. (2009). Entry, exit and labor productivity in uk retailing: evidence from micro data. In *Producer Dynamics: New Evidence from Micro Data*, pages 271–302. University of Chicago Press.
- Lentz, R. and Mortensen, D. T. (2008). An empirical model of growth through product innovation. *Econometrica*, 76(6):1317–1373.
- Maliranta, M. and Määttänen, N. (2015). An augmented static olley–pakes productivity decomposition with entry and exit: Measurement and interpretation. *Economica*, 82:1372–1416.
- Melitz, M. J. and Polanec, S. (2015). Dynamic olley-pakes productivity decomposition with entry and exit. *The Rand Journal of Economics*, 46(2):362–375.
- Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment. *Econometrica*, 64(6):1263–1297.
- Scarpetta, S., Hemmings, P., Tressel, T., and Woo, J. (2002). The role of policy and institutions for productivity and firm dynamics: Evidence from micro and industry data. *OECD Economics Department Working Paper No. 329, Available at SSRN 308680*.
- Van Biesebroeck, J. (2003). Revisiting some productivity debates. Technical report, National Bureau of Economic Research.
- Van Biesebroeck, J. (2005). Exporting raises productivity in sub-saharan african manufacturing firms. *Journal of International Economics*, 67(2):373–391.

Van Biesebroeck, J. (2008). Aggregating and decomposing productivity. *Review of Business and Economics*, 53(2):122–146.

Appendix A. Literature overview

TABLE A1

Application of levels and logs as a labor productivity measure in past decomposition studies

<i>Contribution</i>	<i>Method</i>	<i>Data</i>	<i>Levels/Logs</i>
Griliches and Regev (1995)	GR	Israel Central Bureau of Statistics (CBS) - Industrial Surveys	Levels
Baily et al. (1996)	BHC (modified)	US Census - Manufacturing Sector	Levels
Davis and Haltiwanger (1999)	FHK	US Census - Manufacturing Sector	Logs
Baily et al. (2001)	GR (modified)	US Census - Annual Survey of Manufactures of the Longitudinal Research Database	Levels
Foster et al. (2001)	FHK	US Census - Manufacturing and Services Sector	Logs
Scarpetta et al. (2002)	GR, FHK	Firm-level data from ten OECD countries: United States, Germany, France, Italy, United Kingdom, Canada, Denmark, Finland, Netherlands, and Portugal	Logs
Bernard et al. (2003)	FHK	Simulated data - Based on parameters from US Manufacturing	Levels
Disney et al. (2003)	BHC, GR, FHK	UK Census of Production - Annual Census of Production Respondents Database	Logs
Van Biesebroeck (2003)	BHC (modified)	US - Automobile Assembly Plants and Longitudinal Research Database	Logs
Bartelsman et al. (2004)	GR, FHK	Firm-level data from 24 countries	Logs
Van Biesebroeck (2005)	BHC (modified)	Firm-level data from nine African countries: Based on surveys in the Manufacturing Sector	Logs
Foster et al. (2006)	FHK	US Census - Census of Retail Trade	Logs
Hakkala (2006)	GR, FHK	Statistics Sweden - Sample Manufacturing Sector	Levels
Lentz and Mortensen (2008)	FHK	Danish Business Statistics Register - Annual panel of privately owned firms	Levels
Bartelsman et al. (2009)	GR, FHK	Firm-level data from fourteen countries: Estonia, Hungary, Latvia, Romania, Slovenia, Argentina, Brazil, Chile, Colombia, Mexico, Venezuela, Indonesia, South Korea, and Taiwan [China]	Logs
Haskel and Sadun (2009)	FHK	UK Annual Respondents Database (ARD) - Retail Sector	Logs
Maliranta and Määttänen (2015)	OP (augmented)	Statistics Finland - Structural Business Statistics Data	Logs
Melitz and Polanec (2015)	DOPD	Slovenian AJPES - Slovenian Manufacturing Sector	Logs, (Levels)
Decker et al. (2017)	DOPD, FHK	US Census - Revenue-enhanced Longitudinal Business Database (ReLBD)	Logs
Acemoglu et al. (2018)	FHK	US Census - Manufacturing Sector	Levels
Alon et al. (2018)	DOPD	US Census - Revenue-enhanced Longitudinal Business Database (ReLBD)	Logs
Brown et al. (2018)	DOPD	Mexico - Annual Industrial Survey (EIA) Columbia - Manufacturing Survey (EAM) Chile - National Annual Manufacturing Survey (ENIA) Peru - Annual Economic Survey (EEA)	Levels, (Logs)
Dias and Marques (2021b)	DOPD/FHK (modified)	Statistics Portugal - Portuguese non-financial firms	Logs

Notes: This table provides an overview of recent decomposition literature and documents the measure deployed for representing firm-level productivity (levels and/or logs). BHC: Baily, Hulten and Campbell (1992), GR: Griliches and Regev (1995), FHK: Foster, Haltiwanger and Krizan (2001), DOPD: 'Dynamic Olley-Pakes Decomposition' by Melitz and Polanec (2015); 'Levels' and 'Logs' in parentheses means that some results were reported in these measures as a supplement to the mainly applied measure.

Appendix B. Industry classification

TABLE B1
A38 and ISIC industry classification in manufacturing

A38	ISIC	Description
CA	10-12	Manufacture of food products, beverages and tobacco products
CB	13-15	Manufacture of textiles, wearing apparel, leather and related products
CC	16-18	Manufacture of wood and paper products; printing and reproduction of recorded media
CD	19	Manufacture of coke and refined petroleum products
CE	20	Manufacture of chemicals and chemical products
CF	21	Manufacture of basic pharmaceutical products and pharmaceutical preparations
CG	22-23	Manufacture of rubber and plastics products, and other non-metallic mineral products
CH	24-25	Manufacture of basic metals and fabricated metal products, except machinery and equipment
CI	26	Manufacture of computer, electronic and optical products
CJ	27	Manufacture of electrical equipment
CK	28	Manufacture of machinery and equipment n.e.c.
CL	29-30	Manufacture of transport equipment
CM	31-33	Other manufacturing; repair and installation of machinery and equipment

Notes: The table sets out the intermediate SNA/ISIC aggregation A38, which aggregates similar ISIC two-digit divisions to 13 different categories. It is the industry classification deployed in the main text, excluding the industry of coke and refined petroleum products (ISIC 19).

Appendix C. Summary statistics

TABLE C1
Summary statistics for French manufacturing 2009-18

	Obs	Mean	Sd	Min	Med	Max
All manufacturing						
Value-Added	260,674	5,951,312	$4,81 \cdot 10^7$	1,405.84	1,309,477	$6,04 \cdot 10^9$
Employees (FTE)	260,674	77.14	519.53	5	23	68326.5
Working Hours	260,674	117,363	791,967.40	7,177.97	35,211	$1,04 \cdot 10^8$
Value-Added/Working Hours	260,674	40.64	20.44	0.09	36.30	724.62
Manufacturing of chemicals and chemical products (A38 CE / ISIC 20)						
Value-Added	8,252	$1,47 \cdot 10^7$	$4,42 \cdot 10^7$	65,269.38	3,302,234	$8,96 \cdot 10^8$
Employees (FTE)	8,252	139.16	422.65	5	40	8447.5
Working Hours	8,252	208,831.10	634,371	7,455.88	59,790.08	$1,26 \cdot 10^7$
Value-Added/Working Hours	8,252	61.01	38.49	2.89	52.35	486.37

Notes: The numbers for 'All manufacturing' include firms of all manufacturing industries with the exception of the coke and refined petroleum products industry (ISIC 19). The number of employees is documented in the form of full-time equivalents (FTE). The statistics are reported for the cleaned sample. Value-Added is reported in deflated €.

Appendix D. Log distortions in related decomposition methods

The aim of this section is to show that the general patterns we have identified for the FHK decomposition also hold for the related methods by Griliches and Regev (1995) (GR) and by Melitz and Polanec (2015) (DOPD: 'Dynamic Olley-Pakes Decomposition'), which represent two commonly used alternatives. In addition to the similarities, we will note important differences in the DOPD method.

Like the FHK method, the method by Griliches and Regev (1995) is a longitudinal approach. The GR method decomposes productivity using average weights. For simplicity, we express the decomposition in a somewhat 'neutral' form, not differentiating between levels and logs in the denotation, as done, for instance, by Baily et al. (2001).

$$\begin{aligned} \hat{\Phi}_{GR} = & \sum_{i \in S} \underbrace{\bar{s}_i \cdot \Delta \varphi_i}_{WFE_i} + \sum_{i \in S} \underbrace{\Delta s_i \cdot (\bar{\varphi}_i - \bar{\Phi})}_{BFE_i} \\ & + \sum_{i \in N} \underbrace{s_{i2} \cdot (\varphi_{i2} - \bar{\Phi})}_{N_i} + \sum_{i \in X} \underbrace{s_{i1} \cdot (\bar{\Phi} - \varphi_{i1})}_{X_i} \end{aligned} \quad (D1)$$

As with the FHK method, when using levels, the decomposition formula above would require a reference productivity for calculating growth rates. Since Griliches and Regev use average weights, the choice of $\bar{\Phi}$ may be the most intuitive one in this case (Van Biesebroeck, 2008). However, in line with Baily et al. (2001), who deployed a modified version of the GR method, we use Φ_1 as a reference productivity. This method also facilitates a comparison between our results for the GR and the FHK method. By using averages for firm-level productivity ($\bar{\varphi}_i$) and market shares (\bar{s}_i), there is no interaction term or cross-firm effect, as in the FHK method. A further important difference from the FHK method is the choice of benchmark productivity, which measures the impact of market share reallocations and entering and exiting firms. Instead of the initial aggregate productivity, Φ_1 , the GR method deploys the average between the aggregates in the starting and ending period, $\bar{\Phi}$.

The log distortions for the weighted average industry for the 2009-18 period are shown in Table D1. The results are similar to those presented for the FHK method. The distortions in aggregate productivity growth ranged from -0.64 to 1.36 percentage points and were mostly driven by the distortions in the WFE. The exit and entry components also showed considerable absolute distortions, while the BFE was less affected. Analogous to the results for the FHK method, on average, logs underestimated the WFE and the entry component, but they overestimated the exit component. The BFE showed a minor negative tendency. Taken together, this induces, on average, a positive log distortion in aggregate productivity growth.

The decomposition method by Melitz and Polanec (2015) is based on the cross-

TABLE D1
Decomposition of the log distortions in the GR decomposition

	ε_W	ε_B	ε_X	ε_N	ε_A
2010	-0.22	0.00	-0.59	0.17	-0.64
2011	-0.04	0.11	-0.36	0.17	-0.12
2012	0.21	-0.15	-0.03	-0.06	-0.03
2013	0.02	0.01	-0.29	0.38	0.12
2014	0.21	-0.01	-0.26	0.22	0.16
2015	0.35	-0.01	-0.30	0.25	0.28
2016	0.23	-0.06	-0.18	0.38	0.37
2017	1.10	0.01	0.02	0.22	1.36
2018	-0.10	-0.05	0.07	0.06	-0.02
Mean	0.19	-0.02	-0.21	0.20	0.16
SD	0.39	0.07	0.21	0.14	0.53
Median	0.21	-0.01	-0.26	0.22	0.12

Notes: The table sets out the annual decomposed log distortions during the 2009-18 period for the entire manufacturing sector, using the decomposition method by Griliches and Regev (1995) (GR). The results are based on the annual averages of the industry-level results for the 12 industries in our sample. As industry weights, we used labor input in the form of hours worked, averaged over the beginning and ending years of the period in which the respective growth rate was measured. All reported values for the log distortions are in percentage points.

sectional approach proposed by Olley and Pakes (1996). Instead of tracking individual firms over time, the DOPD method decomposes aggregate productivity in two different periods and subsequently contrasts the individual components. Apart from the entry and exit components, they decompose the contribution of incumbents into a within-firm and a between-firm effect.

$$\begin{aligned}
 \hat{\Phi}_{DOPD} = & \underbrace{\Delta \bar{\varphi}_S}_{WFE} + \underbrace{\Delta cov_S(\varphi_{it}, s_{it})}_{BFE} \\
 & + \underbrace{s_{N2} \cdot (\Phi_{N2} - \Phi_{S2})}_N + \underbrace{s_{X1} \cdot (\Phi_{S1} - \Phi_{X1})}_X
 \end{aligned} \tag{D2}$$

In the DOPD decomposition, the WFE is represented by the development of the unweighted average of firm-level productivity in surviving firms. The BFE is expressed by the change in the covariance between the firm-level productivity of incumbents and their market shares. The last two terms represent the contribution of entering and exiting firms relative to the aggregate productivity of surviving firms at the respective point in time.

When representing firm-level productivity in logs, the above equation can simply be used in the above form. For levels, however, the equation requires a slight modification to ensure scale invariance in the covariance term. Melitz and Polanec (2015) provide a level representation of the decomposition method in the appendix of their paper. Note that they deployed $\bar{\Phi}$ as a reference productivity. For the results presented in Table D2, we followed their suggested approach.

The distortions in aggregate productivity growth ranged from -0.87 to 1.14 percent-

TABLE D2
Decomposition of the log distortions in the DOPD decomposition

	ε_W	ε_B	ε_X	ε_N	ε_A
2010	-0.42	-0.03	-0.62	0.20	-0.87
2011	-0.22	0.10	-0.40	0.19	-0.33
2012	-0.02	-0.05	0.01	-0.19	-0.25
2013	0.08	-0.15	-0.29	0.39	0.02
2014	-0.01	0.08	-0.29	0.23	0.01
2015	0.27	0.03	-0.32	0.25	0.22
2016	0.63	-0.62	-0.19	0.39	0.21
2017	0.73	0.23	-0.05	0.22	1.14
2018	0.33	-0.62	0.11	0.04	-0.14
Mean	0.15	-0.11	-0.23	0.19	0.00
SD	0.38	0.31	0.22	0.18	0.54
Median	0.08	-0.03	-0.29	0.22	0.01

Notes: The table sets out the annual decomposed log distortions during the 2009-18 period for the entire manufacturing sector, using the decomposition method by Melitz and Polanec (2015) (DOPD). The results are based on the annual averages of the industry-level results for the 12 industries in our sample. As industry weights, we used labor input in the form of hours worked, averaged over the beginning and ending years of the period in which the respective growth rate was measured. All reported values for the log distortions are in percentage points.

age points. What stands out as a striking difference between the DOPD method and the FHK and GR methods, is that the BFE is subject to significantly stronger distortions, ranging from -0.62 to 0.23 percentage points with a negative mean. This result implies that, in contrast to the other two methods, whose reallocation component was, in absolute terms, affected only to a limited extent, each of the four productivity components in the DOPD decomposition may be considerably distorted by the use of logs. Moreover, the log distortions in the BFE show, on average, a clearly negative tendency, implying that logs tend to overestimate the BFE component in the DOPD method. This negative tendency also appears to balance out the distortion in aggregate growth, resulting in an average distortion of approximately zero. Hence, for the DOPD method, our sample shows that calculations based on logs are, on average, on spot with respect to aggregate growth.