Papers in Evolutionary Economic Geography

20.20

Structural accounting: an empirical assessment of cross-country differences in productivity

Dario Diodato

Utrecht University Human Geography and Planning

Structural accounting: an empirical assessment of cross-country differences in productivity*

Dario Diodato

dario_diodato@hks.harvard.edu

Center for International Development, Harvard University - 79 JFK street, 02138 Cambridge MA, USA – Current address: Joint Research Centre, European Commission - Calle Inca Garcilaso 3, 41092 Seville, Spain

Abstract

This paper proposes a method to decompose cross-country differences in productivity (TFP) into a technological component - depending on the overall productivity of a country - and an allocation component, which depends on whether factors of productions are allocated to productive or unproductive industries. Using a sample of over 2 million firms from 30 countries, the analysis estimates that 1/4 of inequality between countries is due to the Composition effect, while 3/4 to the Place effect. Moreover, once accounting for heterogeneity at the subnational level, I find that the Composition effect may be as high as 50%.

*Acknowledgments: I acknowledge the financial support of NWO, Innovation Research Incentive Scheme (VIDI). Project number: 45211013. I am truly grateful to Ricardo Hausmann and Frank Neffke, for our lengthy discussions on this paper, to Stefano Breschi for suggestions and for sharing his scraping code in Ruby and to Dani Rodrik for comments. Additionally, I would like to thank the participants of the NBER productivity lunch and CID Growth seminar for their feedback.

JEL classifications: O14, O33, O47.
Keywords: Total factor productivity, factor allocation, structural change.

1 Introduction

Why is it that, for instance, Germany and Japan have high standards of living? Economists typically organize the answer to this central question in two steps: (1) Germany and Japan are wealthier because they are more productive, and (2) they are more productive because of institutions, geography, culture, history (and so forth). Yet, looking closely to the economy of these countries, it is easy to notice that both are specialized in the production of cars, which happens to be one of the most productive industries. Are Germany and Japan more productive in anything they do? Or can it be that Germany and Japan are more productive because they are specialized in industries with higher margins? Moreover, not every part of these countries is as productive as their core manufacturing and service centers. How much regional heterogeneity in productivity creates distortions on they way we interpret the world?

I argue that our understanding of productivity differences is not complete until we unpack what happens at the sub-national level, both in the industry and regional dimensions. The disaggregation of productivity is an additional layer of analysis, a useful tool to better understand the symptoms of economic growth and can, in turn, guide the search for the root causes, as well as help defining growth strategies. The aim of this paper is to shed some light on the matter by measuring the relevance of industrial and regional aggregation in explaining productivity differences. With respect to the industrial dimension, in particular, the purpose is not to causally link industrial structure to wealth, but to estimate the extent to which differences in the output mix are a significant phenomenon. This can, in turn, guide future research into the root causes of growth and, ultimately, help shaping policy strategy.

For early development economists, structural change was a key ingredient of economic growth (Kuznets, 1966). Its role in growth theory became somewhat marginal when the literature focused on models that could explain the Kaldor facts (first neoclassical, then endogenous growth theory), and only recently has the literature begun to reconcile leading models with structural change (Herrendorf et al., 2014). Yet, a number of theoretical and empirical articles in the past two decades have suggested that the composition of a country's output might be important. Models that lead to these conclusions typically involve asymmetric externalities across industries and an economy open to trade (Matsuyama, 1992; Young, 1993; Galor and Mountford, 2006; Hausmann et al., 2007). The main insight of these types of models is that comparative advantages might induce some countries to specialize in technologically stagnant industries, in turn creating differences in the aggregate productivity level or growth. Additional empirical exploration hints in the same direction. Imbs and Wacziarg (2003) find a non-monotonic relationship between the diversification of the output mix and development. Caselli (2005) finds that structural change – in this case shifting factors of production to non-agricultural sectors – would have nearly the same effect on reducing cross-country income variance that one would obtain by increasing agricultural productivity. This is because productivity in agriculture exhibits much greater variance than does productivity in the rest of the economy. A similar point is made in Rodrik (2013), showing that manufacturing exhibits unconditional convergence. Rodrik (2013) concludes that the lack of convergence observed at the national level cannot be due to economy-wide factors.

Yet, in spite of this emerging literature we still do not have a clear expectation on how much the variations in output mix contribute generating differences in overall productivity. Since this has, as argued, important implications on how we study the causes of development and how we shape growth policy, in this paper I measure the relevance of output structure in explaining productivity differences. To this end, I explore in detail total factor productivity (TFP) and its industrial dimension, decomposing the variance of TFP into two main components. The first, a Composition effect, measures the part of the variance that depends on the industrial structure of a nation. Nations with high shares in productive industries would have a strong Composition effect. The second, a Place effect, captures the amount of variance that can be attributed to nationwide differences in TFP. When a nation, for instance, is relatively more productive in all industries, the Place effect plays a more important role.

I estimate TFP, the Place effect, the Composition effect and their variance using a firm-level dataset of firms spread across several nations and industries (Bureau van Dijk[1]). The cross-country estimates from this analysis reveal that the Place effect explains approximately 79% of the difference, while 29% can be attributed to the Composition effect.[2] This implies that, even if it were possible to freely transfer technology across countries, large differences in standards of living would persist. It is argued that this cross-country estimate of the Composition effect is a lower bound for two reasons. First, adjusting for price differences across countries indicate that the Composition effect might be more important. Second, countries are collections of heterogeneous cities and regions, some more wealthy, urbanized and diversified, others more rural typically lagging in productivity and lacking diversity. Accounting for subnational heterogeneity suggests a much larger role of Composition of up to 50%.

A case is made that the relationship between the Composition and Place effects can be additionally interpreted as an indicator of how easily technology diffuses through space: the more dominant is the Composition effect, the more places are technologically homogeneous (that is, with lower variance of within-industry productivity). This intuition is exploited to run the decomposition analysis on different sets of nations and regions, with the expectations

[1]To compromise between breadth and depth, two datasets from Bureau van Dijk are used: the main dataset is a cross-section from Orbis. For particular estimates, I use a 10-year (2005-2014) panel of European firms from Amadeus, which also includes information about the NUTS3 location of firms.

[2]Composition and Place do not sum to one. The next section details the mathematical properties and possible interpretations of the measure.

that in more homogeneous places (proximity, language, integrated markets) the Composition effect plays a more relevant role. In line with this expectation, I find that subnationally, for instance in France, Germany and Italy, the Composition and Place effects are responsible for (roughly) an equal amount of inequality across regions.

The main message of the paper is that the industrial composition of countries explains an important share of the differences in standards of living. This allows for old and new explanations that – complementary to the standard approach – can improve our understanding of growth and development.

The results of the analysis can be related to other analyses of structural change. Articles such as Timmer and de Vries (2009) or McMillan et al. (2014) find that structural change in recent years, especially in Latin America and Africa, has been growth reducing (with the trend reversing in Africa after 2000). Whereas this type of analysis is concerned with change over time (the amount of labor that moves from and to productive and unproductive sectors), this paper offers a cross-sectional picture, to clarify how large the potential gain from structural change and technological catch up is. This method is more similar to that in Caselli (2005), which assesses the gains from changes in agricultural productivity *vis-à-vis* the gains from transitioning to non-agricultural activities. This analysis expands this methodology and applies it to a dataset that includes as many as 800 sectors. Another related stream of literature is emerging around the so-called misallocation hypothesis (Banerjee and Duflo, 2005). The conjecture is that a large portion of the within-industry productivity gap between developed and emerging economies is due to the latter having high productivity dispersion. The lower aggregate sectoral productivity is the result of factors of production being disproportionately employed in unproductive firms. This hypothesis finds some empirical confirmation (Hsieh and Klenow, 2009 and 2010; Bartelsman et al., 2013), and it has been interpreted as the consequence of a distortion (e.g., different rules for small firms) in the allocation and selection process. Comparably, this paper examines the role of allocation and selection, not within but between industries.

The paper proceeds as follows. Section 2 presents a novel method of cross-sectional productivity decomposition. Section 3 discusses the data, as well as issues related to the measurement of productivity. In section 4, one can find stylized facts about productivity. The main results of the paper are in section 5, while section 6 concludes the paper.

2 Decomposing productivity: Composition effect and Place effect

Using standard notation, I define TFP as

$$A_{ig} = \frac{Y_{ig}}{K_{ig}^{\alpha} L_{ig}^{1-\alpha}}, \tag{1}$$

where g is used to index locations (be it countries or regions) and i to index industries. Equation 1 is a standard measure of TFP. Y_{ig} is output, K_{ig} is capital, and L_{ig} is employment. This definition of productivity serves as a reference. In this paper, a variety of productivity measures are explored and tested for robustness. For instance, I allow for capital and labor shares to be estimated, output to be measured as revenue or quantity and employment to be adjusted by wages to control for differences in human capital. The next section elaborates on the measurement of TFP. For the present, the definition in 1 is sufficient for our purposes. The share of industry i in country g is defined as

$$s_{ig} = \frac{K_{ig}^{\alpha} L_{ig}^{1-\alpha}}{K_{g}^{\alpha} L_{g}^{1-\alpha}}. \tag{2}$$

Then, aggregate productivity can be expressed as follows

$$A_{g} = \frac{Y_{g}}{K_{g}^{\alpha} L_{g}^{1-\alpha}} = \sum_{i} s_{ig} A_{ig}. \tag{3}$$

The decomposition measure is derived from the following identity.

$$\sum_i (A_{ig} - A_i)(s_{ig} - s_i) = \sum_i (A_{ig} - A_i)(s_{ig} - s_i), \tag{4}$$

where s_i is the share of industry i in the whole economy and is defined similarly to s_{ig} in equation 2. The identity can be rewritten as

$$A_g = -A + \underbrace{\sum_i s_{ig} A_i}_{\text{Composition}} + \underbrace{\sum_i s_i A_{ig}}_{\text{Place}} + \underbrace{\sum_i (A_{ig} - A_i)(s_{ig} - s_i)}_{\text{Residual}}. \tag{5}$$

The Composition term, which this paper also refers to as C_g, indicates the amount of TFP that can be attributed to the output structure. The Place term, also called P_g, captures the amount of TFP due to own productivity – regardless of the output mix. The third term, R_g, can be seen as a residual, but it could also have an economic interpretation as a Ricardian effect of specialization: R_g is high if a place g has large shares in activities in which it has high productivity (although the advantage in R_g is absolute, not relative). As is shown below, this residual/Ricardian term plays no role in the main indicator employed in this analysis, and therefore, the first interpretation, as a residual, is stressed. However a secondary indicator used for robustness is included in the analysis. Note that the term A is the average productivity of the economy, and since it is constant, it does not contribute to the variance. We have

$$Var(A_g) := \sigma_A^2 = i' \begin{pmatrix} \sigma_C^2 & \sigma_{CP} & \sigma_{CR} \\ \sigma_{CP} & \sigma_P^2 & \sigma_{PR} \\ \sigma_{CR} & \sigma_{PR} & \sigma_R^2 \end{pmatrix} i. \tag{6}$$

Where i is a vector of ones, and the matrix on the right-hand side (RHS) is the variance-covariance matrix of Composition (C), Place (P), and the Ricardian/residual effect (R). To obtain the main indicators, the following thought experiment is employed. What would the variance of productivity be if every country had the same industry productivity? If we set $A_{ig} = A_i$, the standard deviation σ_A becomes equal to σ_C. Equally, we can ask, what would the variance of productivity be if every country had the same composition of output? Will this reallocation of production factors reduce the variance? We can see by substitution that, by setting $s_{ig} = s_i$, we simply obtain that $\sigma_A = \sigma_P$. Our main indicators are

$$
\begin{aligned}
I_P &= \frac{\sigma_P}{\sigma_A}, \\
I_C &= \frac{\sigma_C}{\sigma_A}.
\end{aligned} \tag{7}
$$

The indicator I_P describes the world, at current productivity, with no Composition effect. It reports what share of the standard deviation of productivity would be left if we were to shut down the Composition effect entirely and is, therefore, an indicator of the importance of the Place effect. Symmetrically, I_C indicates what share of the standard deviation of productivity would be left by nullifying productivity differences. It is, then, an indicator of the Composition effect. Additional interpretation can be unveiled by noting that

$$Var(P_g) \coloneqq \sigma_P^2 = Var(\sum_i s_i A_{ig}) = \sum_{i,j} s_i s_j Cov(A_{ig}, A_{jg}). \tag{8}$$

The covariance term is high when productivity is mainly a country characteristic: for any industry pair ij, higher productivity in country g in i implies higher productivity in j. The effect is amplified by the weights, when the covariance is high in industries that represent a large share of the economy. When the covariance term nears zero for all industry pairs, there is no Place effect and industry productivity is randomly distributed across countries. For the Composition effect, we can write

$$Var(C_g) \coloneqq \sigma_C^2 = Var(\sum_i s_{ig} A_i) = \sum_{i,j} A_i A_j Cov(s_{ig}, s_{jg}). \tag{9}$$

The Composition effect is irrelevant when industry shares are randomly distributed. It becomes more important for explaining the variance of productivity across countries when some places have systematically larger shares in industries that are more productive on average.

An interesting link with equation 9 is that the term $Cov(s_{ig}, s_{jg})$ is one of the common metrics to measure coagglomeration economies or the product space (Porter, 2003; Hausmann et al., 2014; Diodato, Neffke, and O'Clery, 2018b). The idea behind these types of indicators is that two industries co-occur geographically more often when they have specific bilateral externalities, such as Marshallian knowledge spillovers, labor sharing and supply-chain linkages (Ellison et al., 2010). As discussed in the introduction – theoretical models, which support the hypothesis that the output mix is important to understand productivity differences across countries, typically resort to industry-specific externalities. It is interesting, then, that the importance of composition in the variance depends on the degree to which industries coagglomerate.

3 Data and measurement of productivity

To measure country-industry productivity, I start from firm-level statistics from Bureau van Dijk's Orbis dataset, which collects balance sheet information for more the 200 million companies. The main shortcoming of this data is that there is some variation in terms of coverage between countries. Orbis includes all companies that are legally obliged to publicly report their accounts. Differences in laws that trigger mandatory reporting in different countries create heterogeneity in coverage. Bureau van Dijk attempts to correct this bias by collecting information on all the remaining companies that they can. This is unlikely to entirely solve the problem, but it certainly mitigates the distortion. I show, in fact, in section 5.1 that firms in the Orbis dataset represents reasonably their countries' economy. Orbis has many other advantages. It is a global dataset, which allows for systematic comparison across the world. Most companies are assigned detailed geographical (including subnational information) and industrial codes,[3] which serve the purpose of this research. One can also distinguish consolidated from unconsolidated accounts, to avoid assigning attributes to a geographical area that in fact come from the accounts of a multi-plant (potentially multinational) corporation. Finally, Orbis contains balance sheet data that can be used to measure TFP.

In its simplest form, TFP can be measured using value added (VA) as output, fixed capital as capital stock K and the number of employees as labor L. However, recent works (see, for instance, Bartelsman et al., 2013) have stressed the difference between TFPR (TFP revenue) and TFPQ (TFP quantity). It can be easily seen with a stylized model that TFPR is not an adequate measure of firm productivity. If we describe the environment of firm j with a model of monopolistic competition and a technology that uses capital and labor as input, demand for firm j is

$$x_j = p_j^{-\sigma} \frac{E}{I}, \tag{10}$$

where x is physical output, p is price, σ is the elasticity of substitution,[4] and E and I are total demand and the price index of the economy,[5] respectively. Production is

$$x_j = A_j K_j^{\alpha} L_j^{1-\alpha}, \tag{11}$$

where A_j is TFP. However, measuring productivity as TFPR would result in

$$TFPR_j = \frac{VA_j}{K_j^{\alpha} L_j^{1-\alpha}} = \frac{p_j x_j}{K_j^{\alpha} L_j^{1-\alpha}} = A_j p_j = \frac{\sigma}{\sigma - 1}(\frac{r}{\alpha})^{\alpha}(\frac{w}{1-\alpha})^{1-\alpha}. \tag{12}$$

[3]To get a sense on the level of detail, the dataset distinguishes between 16 German NUTS1 (Länder), 38 NUTS2 (regions), and over 400 NUTS3 (districts). To get a sense of the industrial detail, transportation equipment is divided in 2 NACE rev2 2-digit industries and 12 4-digit industries, including manufacturing of bicycles, motorcycles, motor vehicles, semi-trailers, ships, sporting boats, trains, airplanes.

[4]Beware of the possible confusion between elasticity of substitution (σ) and standard deviation/variance/covariance (e.g., σ_A, σ_C^2, σ_{PR}).

[5]To lighten notation, I define $I = \sum_k p_k^{1-\sigma}$.

The price of labor and capital is w and r, respectively. To obtain the final expression for $TFPR_j$, two substitutions are applied. First, equation 11 is used to substitute for $K_j^{\alpha}L_j^{1-\alpha}$. Then, optimal pricing behavior would lead to the mark-up rule where $p_j = [\frac{\sigma}{\sigma-1}(\frac{r}{\alpha})^{\alpha}(\frac{w}{1-\alpha})^{1-\alpha}]/A_j$. The term A_j – the objective of the analysis – drops out. This framework shows that there is no theoretical proportionality between a firm's $TFPR_j$ and A_j, the efficiency with which inputs are transformed into output. Intuitively, this lack of correlation stems from the opposing effects of optimal output versus optimal pricing. While output is proportional to productivity, it is the opposite for prices in the event that firms choose a mark-up on marginal costs, which in turn are negatively related to productivity.

Hsieh and Klenow (2009) highlight that there is a theoretical link between output and value added – in this model, it can be derived from equation 10 as $x_j = (p_j x_j)^{\sigma/(\sigma-1)}(I/E)^{1/(\sigma-1)}$ – and propose the following adjusted measure of TFPQ

$$TFPQ_j = \frac{VA_j^{\sigma/\sigma-1}}{K_j^{\alpha}L_j^{1-\alpha}}\kappa. \tag{13}$$

The Hsieh and Klenow (HK) measure is the baseline indicator of productivity in this analysis. In their work, the TFPQ indicator is calculated using the standard capital share of one-third ($\alpha = 0.33$) and an elasticity of substitution $\sigma = 3$. In the benchmark results in this paper, I opt for the same capital share but a larger elasticity ($\sigma = 7$). In the results and in the appendix, one can find a careful explanation for why this benchmark was chosen, as well as a discussion of the robustness of this choice.

I highlight here that using TFPQ has a drawback in this context. Since we wish to aggregate productivity across industries, the aggregation in equation 3 requires that the units of measure for output Y_{ig} are comparable across industries. When Y_{ig} is in values, comparability is guaranteed. However, when Y_{ig} is in output, one would need to resort to the additional assumption of homogeneous units of output for the different industries. For this reason, I also report and discuss TFPR. It should be noted, additionally, that the measure of this paper does not actually use physical quantities, but rather deflate values of final output – rendering the aggregation conceptually less problematic.

An important control in this context concerns the price index. In the HK measure, the term $\kappa = (I/E)^{1/(\sigma-1)}$ is dropped because the authors are interested in comparing the productivity of firms in the same country and in the same industry. They subsequently assume that κ is a constant.[6] This assumption is more stringent in our context, where the comparison between countries and industries is at the core of the analysis. I test whether the main conclusions of the baseline indicator are robust to a country-specific price index, in an attempt to factor in the evident nominal differences across countries. To coherently model a country-specific price index (avoiding the complications of including trade costs), countries must be thought of in isolation. Using the simplifying assumption that each firm j in industry i produces with efficiency A_i,[7] the price index I_g is

$$I_g = \sum_i P_{ig}^{1-\sigma} = [\frac{\sigma}{\sigma-1}(\frac{r_g}{\alpha})^{\alpha}(\frac{w_g}{1-\alpha})^{1-\alpha}]^{1-\sigma}\sum_i(1/A_i)^{1-\sigma}. \tag{14}$$

Demand, E_g, equals national value added, that is $r_gK_g + w_gL_g$. The adjustment term, $(I/E)^{1/(\sigma-1)}$, is therefore

$$(I_g/E_g)^{1/(\sigma-1)} = \frac{r_g^{\alpha}w_g^{1-\alpha}}{(r_gK_g + w_gL_g)^{1/(\sigma-1)}}\kappa_1, \tag{15}$$

where κ_1 collects all the constant terms. The only additional assumption required to construct the adjustment term concerns the rental price of capital, r, which I do not observe. Following HK's reasoning, I set an homogeneous price of 10%, which includes a 5% interest rate and 5% depreciation. Alternative

[6] The term κ in Hsieh and Klenow (2009) also includes a distortion that is not modeled here. This difference does not change the constant nature of κ.

[7] With these assumptions, the correction term accounts only for exogenous differences such as capital and labor endowment or institutional factors, not for productivity differences.

prices of capital are also tested, but there are minimal differences in the outcome (available upon request). While it certainly is an impure way to correct for price differences across countries, the term in 15 behaves as expected – with a negative relationship with both countries' output and productivity.

A last issue in the measurement of productivity concerns the share of capital, α. In the core of the analysis in this paper, I assign to α the widely accepted value – at least for the whole economy – of one-third. Moreover it assigns the remaining two-thirds to labor $(1-\alpha)$. For robustness, the analysis is repeated using estimated values of $\hat{\alpha}$. The assumption of constant returns to scale is also relaxed, using $\hat{\beta}$ instead of $1-\hat{\alpha}$. The use of estimated parameters raises additional concerns. As highlighted in previous works, most notably by Olley and Pakes (1996) and Levinsohn and Petrin (2003), OLS or fixed effects estimates of the production function would lead to biased results. Assuming the same production function as I have in this paper – $Y_{ig} = A_{ig}K_{ig}^{\alpha}L_{ig}^{\beta}$ – the simplest translation into an econometric model would be

$$log(Y_{ig}) = \alpha log(K_{ig}) + \beta log(L_{ig}) + \gamma + \epsilon_{ig}, \tag{16}$$

where A_{ig} would be equal to $exp(\gamma + \epsilon_{ig})$. The issue arises if unobserved shocks in the production technology (ϵ_{ig}) are correlated with K and L, which would occur in the event that firms adjust the level of inputs to the change in productivity. Levinsohn and Petrin (2003) show that the direction of the bias with more than one factor of production is nontrivial. In the simple case in which capital is fixed (i.e., firms cannot adjust the number of machines in a given year in response to an unexpected shock) and labor is a free to move, then α would be unbiased while β would be upward biased. To solve the simultaneity bias, Levinsohn and Petrin (2003) (building on the work of Olley and Pakes, 1996) propose an estimation method – LP hereafter – in which intermediate inputs are used as a proxy variable for the unobserved shock. The results of this paper are tested for robustness to simultaneity bias. A 10-year panel dataset derived from a subsection of Orbis (Amadeus) comprising only European firms is used in this case.

4 Stylized Facts

Productivity at the country or industry level is measured starting from firm-level data. A direct consequence of this choice is that the quality of the estimates for each nation varies with the coverage that Bureau van Dijk has of that country. As discussed in section 3, both the breadth and depth of coverage is rather heterogeneous across countries in this dataset. Moreover, an additional constraint is imposed by the availability of the information required to build a TFP measure, namely employment, fixed assets, and value added. These strong requirements all add up to vastly reduce the pool of firms in the Orbis dataset that we can use. For instance, the data include almost 600,000 Dutch companies. Yet, value added information is available for only 100 of these companies. Since the reliability of the TFP estimate at the country level depends on how many firms are available for that country, I am forced to drop a considerable number of countries from the analysis for which the number of firms was deemed insufficient. Using a threshold of 500 firms, we are left with 30 countries.[8]

Table 1 lists the top-10 and bottom-10 countries by productivity.[9] It can be observed that the majority are developed or emerging economies. While this somewhat limits the extent to which we can generalize this study, it is worth noting that the restricted sample still presents large variance to be explained, with almost 10-fold differences in TFP from the most- to the least-productive country.

[8]These countries are Australia, Austria, Belgium, Bosnia and Herzegovina, Bulgaria, the Cayman Islands, the Czech Republic, Finland, France, Germany, Hong Kong, Hungary, India, Ireland, Italy, Japan, South Korea, Latvia, Luxembourg, Macedonia, Norway, Poland, Portugal, Romania, Slovakia, Slovenia, Spain, Sweden, Ukraine, and the United Kingdom. For robustness, the analysis is also repeated with different thresholds.

[9]While I chose a threshold of 500 firms for the benchmark results, I opted to be more inclusive in this descriptive section and to show countries with information on more than 300 firms. This includes countries like Switzerland and the US.

Table 1: Countries ranked by productivity.

ISO2	Name	A_g	# 4-d sectors	# firms
		Top 10		
CH	Switzerland	0.182	136	322
DE	Germany	0.161	761	73309
JP	Japan	0.155	274	4098
SE	Sweden	0.141	567	90415
AU	Australia	0.139	223	908
FR	France	0.135	579	169024
IE	Ireland	0.128	299	1942
AT	Austria	0.124	501	4849
GB	United Kingdom	0.121	580	34287
FI	Finland	0.121	520	24733
		Bottom 10		
PT	Portugal	0.043	577	217464
CZ	Czech Republic	0.043	683	43886
VN	Viet Nam	0.041	146	453
KR	Korea, Republic of	0.040	406	48692
SK	Slovakia	0.037	539	35402
RO	Romania	0.025	567	79816
BA	Bosnia and Herzegovina	0.020	369	3557
BG	Bulgaria	0.018	569	63888
MK	Macedonia, the Former Yugoslav Republic of	0.018	486	13601
UA	Ukraine	0.018	406	6982

Productivity is measured as the benchmark TFPQ, that is, adjusting value added using $\sigma = 7$ and using fixed assets for capital and employment for labor; the capital share is set to 0.33, and countries with fewer than 300 firms are excluded.

As highlighted in the introduction, it is natural to think of these differences across countries as a result of generalized productivity differences in all sectors. Countries such as Japan or Germany are more productive than Macedonia or Ukraine in most industries. Figure 1 in fact shows that – if we take the 30 least-productive industries – the countries with the largest TFP are still more productive than the average in these low-productivity industries. One can think of figure 1 as a visual depiction of the Place effect: some countries are more productive in everything they do, and hence, data on their general level of productivity provide information about productivity in some particular industry. That is, $Cov(A_{ig}, A_{jg}) > 0$, which means that $Var(P_g) > 0$ (see equation 8).

Figure 1: A visual depiction of the Place effect. Countries with higher productivity are also more efficient in low-productivity industries.

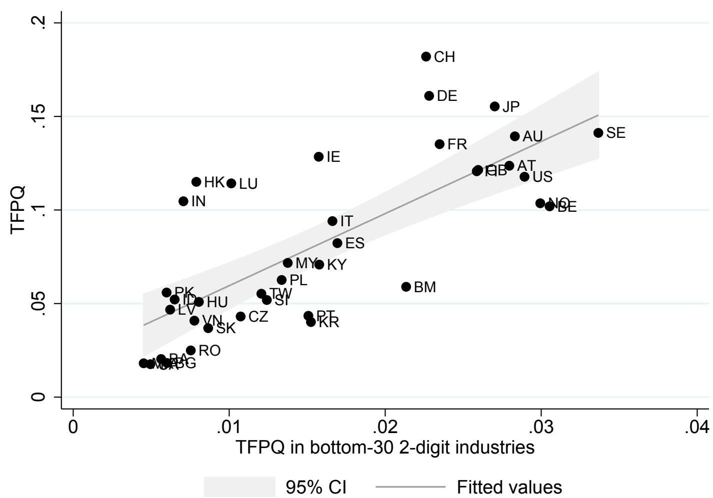

Productivity is measured as the benchmark TFPQ.

However, the average differences in productivity between industries are as large as those across countries. Table 2 in fact shows that telecommunications firms are on average 10 times more efficient at transforming inputs into output than agricultural firms. These large differences in productivity must be considered in combination with the specialization of countries. In fact, one can notice, for instance, how manufacturing of motor vehicles is among the most productive industries. This is an industry in which Germany is traditionally strong and employs large shares of the workforce. This implies that at least part of the higher productivity of this country is due to a specialization in high-productivity industries. Symmetrically, it is well known that in less-developed countries, a larger portion of the population is employed in agricultural activities, hence reducing the average productivity of the country.

Table 2: Industries ranked by productivity.

NACE	Name	A_i	# countries	# firms
		Top 10		
61	Telecommunications	0.238	38	5322
12	Manufacture of tobacco products	0.226	21	121
19	Manufacture of coke and refined petroleu	0.224	32	648
51	Air transport	0.222	34	642
60	Programming and broadcasting activities	0.194	35	1952
9	Mining support service activities	0.194	33	497
35	Electricity, gas, steam and air conditio	0.189	38	8509
21	Manufacture of basic pharmaceutical prod	0.184	37	2006
30	Manufacture of other transport equipment	0.184	36	3241
29	Manufacture of motor vehicles, trailers	0.182	37	6965
		Bottom 10		
75	Veterinary activities	0.067	24	3071
15	Manufacture of leather and related produ	0.064	33	9133
31	Manufacture of furniture	0.064	36	14234
87	Residential care activities	0.060	26	8190
16	Manufacture of wood and of products of w	0.060	35	16593
85	Education	0.058	33	23207
91	Libraries, archives, museums and other c	0.057	21	959
13	Manufacture of textiles	0.042	36	10909
2	Forestry and logging	0.037	30	6753
1	Crop and animal production, hunting and	0.028	38	47940

Productivity is measured as the benchmark TFPQ.

Figure 2: A visual depiction of the Composition effect. Countries with higher productivity have larger shares in high-productivity industries.

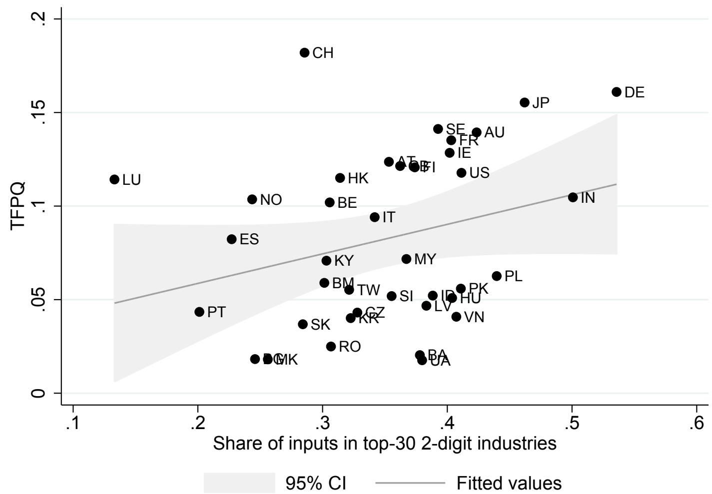

Productivity is measured as the benchmark TFPQ. Shares are computed as in equation 2.

Figure 2 visually captures the concept of the Composition effect. The larger the amount of resources a country devotes to production, in industries that are on average more productive, the larger its average TFP. Together, figures 1 and 2 suggest that, to explain the differences in productivity across countries, both the Composition and the Place effect are required. Starting with the following section, the analysis of this paper, based on the decomposition method proposed in section 2, is aimed at more rigorously evaluating the relative importance of the two factors.

5 Results of decomposition

5.1 The industry dimension

In this section, I present the main results of decomposing productivity differences between countries. That is I use the measure developed in section

2 to assess how much of productivity is due to industries composition. For now, I ignore that we can see countries as collection of regions. The results of decomposition accounting for within-nation regional differences are discussed in section 5.2.

The benchmark estimation in table 3 can be taken as a representative summary of the results. Recall that, for this benchmark, I chose an elasticity of substitution of $\sigma = 7$ and a constant capital share of $\alpha = 0.33$. Moreover, employment instead of the wage bill is used for the labor variables, the economy is divided into 4-digit industries, and all countries with fewer than 500 firms have been excluded from the analysis. With these choices, I find that – conservatively – at least 1/4 of the cross-country inequality in productivity is due to the Composition effect. More precisely, this means that, if we were to erase productivity differences within industries overnight, approximately 29% of the standard deviation of productivity would persist due to specialization in industries with different productivity levels. Moreover, over 3/4 of the inequality is due to the Place effect. More accurately, if we were to change the industry share such that every country dedicated an equal amount of resources (equal to the world average) to every industry, 79% of the cross-country variation would remain.

This is an interesting result. While the predominance of the Place effect is probably the prior of most researchers and development practitioners, I show here that in our most conservative estimates – without price corrections and without accounting for within-nation regional heterogeneity, at least 29% of the variation is due to the Composition effect, a number far from negligible. This shows that lagging nations could make progress not only by adopting the technology of more-advanced countries but also by shifting resources within their industrial/technological portfolio.

Table 3: Summary of results.

Estimate type	Composition	Place	Residual
Benchmark	0.292	0.786	
Alternative Indicator	0.193	0.624	0.183
Firm threshold = 250	0.305	0.782	
Firm threshold = 1000	0.299	0.811	
Constant firms (5000)	0.490	0.795	
Unconsolidated accounts	0.235	0.811	
Exclude selected services	0.292	0.786	
Eurostat shares (1-digit industries)	0.160	0.998	
Orbis shares (1-digit industries)	0.163	0.966	
Wage bill instead of employment	0.378	0.573	
Price Index Adjustment	0.476	0.584	
$\sigma = 2$	0.333	0.721	
$\sigma = 12$ ($\backsim$ TFPR)	0.223	0.883	
Estimated K and L shares (OLS)	0.387	0.796	
Estimated K and L shares (FE)	0.541	0.885	
Estimated K and L shares (LP)	0.728	0.932	

Composition and Place effect, with the exception of the alternative indicator in the second row, are expressed in terms of relative standard deviation, that is, σ_C/σ_A and σ_P/σ_A

In the remainder of this section, I provide an in-depth discussion of the robustness of the estimates of the Composition and Place effects. The main table of results – table 3 – reports selected output. An overview of the issues related to robustness is provided here. For some of the issues the main text provides an in-depth discussion, while for others dedicated sections can be found in the appendix. Specifically, in this section I discuss whether the firms selected in the Orbis dataset are representative of their countries' economy and the role of nominal price differences across countries. The sensitivity to the choice of elasticity is discussed in appendix A. The choice of capital share α is discussed in appendix B. Sections 5.2, 5.3, and appendix C discuss changes in the industrial and geographical dimensions.

It is important to highlight that, using the chosen indicator, the Composition and the Place effects do not sum up to one. The variance of productivity is not only generated by the Composition and Place effects, but also by the variance of a third residual term, as well as all the covariances. The choice of solely examining σ_C and σ_P is justified by the fact that setting $A_{ig} = A_i$ yields $\sigma_A = \sigma_C$. Similarly, setting $s_{ig} = s_i$, $\sigma_A = \sigma_P$. These thought experiments (what would σ_A be in the two cases) allow us to organize our understanding of the structure of A in two parts: Composition and Place. Since this is already adding a rich layer of complexity to the way normally the structure of A is understood, I opted for this bipartite division. However, a valid alternative would have been to analyze all terms that compose the variance of A_g. In fact, the following equivalence also holds

$$Var(A_g) = Cov(P_g, A_g) + Cov(C_g, A_g) + Cov(R_g, A_g) \tag{17}$$

where, for instance, $Cov(P_g, A_g) = Var(P_g) + Cov(P_g, C_g) + Cov(P_g, R_g)$. In the second row of table 3, this property is exploited to construct a second indicator, where the three terms sum up to one. It is clear that the importance of the Composition and Place effects (at least in relative terms) is not excessively different from that estimated with the primary indicator, with the former capturing 20% and the latter 60%. This alternative, however, leaves room for a residual term. As discussed in section 2, the residual term can also be interpreted as Ricardian specialization. According to this interpretation, approximately 20% of the differences in productivity can be ascribed to countries specializing in what they are better at.

Table 3 also shows a number of robustness checks aimed at addressing the concern that Orbis might not be representative. First, a crucial check is whether the presence of firms with consolidated accounts are significantly distorting the results, but it does not seem to be the case. Next, a small threshold in terms of the minimum number of firms required for a country to enter the analysis is concerning because we might include a country where the information collected in Orbis are not representative of its structure. I observe however that results are robust to different thresholds, with minimal changes halving or doubling the threshold. This means that the results are not driven by the inclusion/exclusion of few extra countries where information are more likely to be biased. Next, in the row labeled 'constant firm', I include only the 19 countries that have more than 5000 firms represented. I then select, randomly, 5000 firms for each of the countries, to keep the sample size constant. This process, which if repeated gives consistent results, suggests that the potential misrepresentation of some countries does not bias the estimate of the Place effect, but it might bias the Composition effect. This hints at a potentially larger role of the Composition effect. We can additionally check whether the firms in Orbis are over- or under-represented in some industries. To control whether that is the case, I use value added by 1-digit industry (21 NACE aggregated industries, identified by a letter code) as reported in the national accounts of individual countries (use and make tables). I collect the national accounts for 18 European nations (Eurostat) and for each of them compare the share of value added of every 1-digit industry as computed using Eurostat with the same share obtained using the Orbis dataset.

Figure 3 shows this comparison for Germany, whose pattern is very similar to that of all other 17 nations checked (with the exception perhaps of Spain, which has slightly larger discrepancies[10]). The figure (left panel) shows that some industries, particularly public services (such as public administration, code O), but also some privately supplied services real estate and professional and scientific activities (L and M) have significant distortions. Code C contains all manufacturing activities and also shows non negligible bias. The right panel of figure 3 shows, nonetheless, that – once we exclude all public services, as

[10]Versions of figure 3 using the 17 other nations are available upon request.

well as industries L and M – there is a better correspondence between the shares estimated with the national accounts and those estimated with Orbis data. Manufacturing too appears to have better estimated shares in this case. I interpret this pattern, which again is not only observed in Germany, as evidence that Orbis systematically captures some industries better than others.

Figure 3: Distribution of activities across industries in Germany according to Orbis and Eurostat.

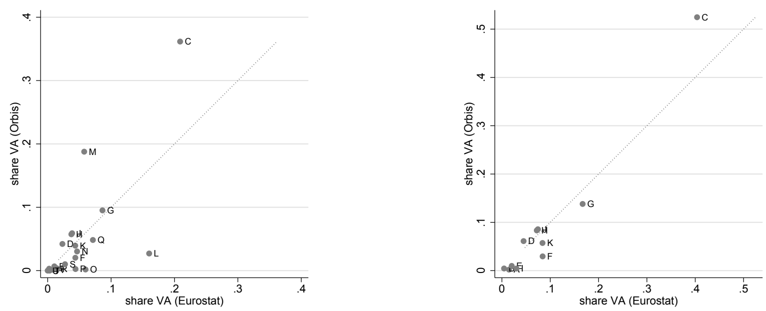

The left panel includes the whole economy, the right panel shows only selected industries.

To verify whether this could be a problem I perform two controls: first, I test the decomposition excluding all the industries that show here to be poorly represented in Orbis. In table 3 the row ‘exclude selected services’ shows that outcome of a decomposition where all industries with 4-digit NACE code larger than 6700 have been dropped.[11] Only marginal changes can be observed. Second, I attempt to use the industry shares estimated with the national accounts directly in the decomposition. While in later discussion in this paper I argue that using coarse industrial aggregation results in biased outcome, I can run here two decomposition routines, one using Eurostat shares the other using Orbis shares for 1-digit industries. The results, because of the coarse aggregation (all manufacturing is collapsed in one sector) are not very informative alone (besides, it also uses a smaller selection of countries). However, comparing one against the other suggests that Orbis is representative at 1-digit level

[11]Mainly: real estate, professional services, public administration, defense, education, health and social services, arts and entertainment.

for our purposes. While we cannot test the same with more disaggregated industrial classifications, the results are reassuring.

Another important set of controls is performed on the sensitivity of results to price corrections. Table 3 shows three types of decomposition in its third block. The first one uses employees payroll instead of employment as in Hsieh and Klenow (2009), following the argument that wages correct (imperfectly) for differences in human capital. The second one uses the correction term theoretically derived in section 3. For both estimates, the role of Composition is closer to that of the Place effect, with 38% Composition and 57% Place in the first case, and 48% Composition and 58% Place in the second one. For this reason, I regard the benchmark estimate of 1/4 Composition and 3/4 Place as a lower bound.

Next, the estimate of the Composition effect increases at lower elasticities of substitution, σ, and declines when the elasticity raises. The outcome of the decomposition is, however, relatively stable around the benchmark. As a midpoint between the values found in the literature of $\sigma = 3$ and $\sigma = 12$ (which gives an approximation of TFPR, as in this case $\sigma/(\sigma - 1) \simeq 1$), the value of $\sigma = 7$ is chosen as a compromise because neither option is desirable. Appendix A elaborates on this issue.

The capital and labor shares (α and β) estimated via OLS do not substantially change the outcome of the decomposition, with only the Composition effect increasing to 39%, while the Place effect remains stable. However, this is not the case when estimating the shares using fixed effects (FE) or the Levinsohn and Petrin estimator (LP). In this case – see appendix B – the estimated shares of capital and labor show a net departure from the assumption of constant returns to scale. The resulting decomposition shows a marked increase in both the Composition effect and Place effect. While such large differences are likely to reflect issues with LP estimates (LP was expected to be more in accord with OLS than FE, see Levinsohn and Petrin, 2003), this shows that the decomposition of productivity in this dataset is not robust to the estimator used. Nonetheless, the qualitative conclusion that both effects are important, with the Place effect capturing the largest share of the variance, continues to hold.

5.2 Aggregation and the regional dimension

An important question concerns the role of geographical heterogeneity into the estimates. It is a well-established fact that levels of productivity can exhibit large differences across regions in a country. Using OECD data, for instance, I verify that the ratio between richest and poorest NUTS2 region in a country (based on income per capita PPP in 2013) is 1.7 for France and Germany, 1.8 for the UK, 2 for Spain, and 2.2 for Italy. The ratio between the richest region in Italy and the poorest in Germany is 1.5. This heterogeneity suggests that it is not always meaningful to think about the difference in productivity between two countries.

Figure 4: Aggregation in the geographical and industrial dimensions.

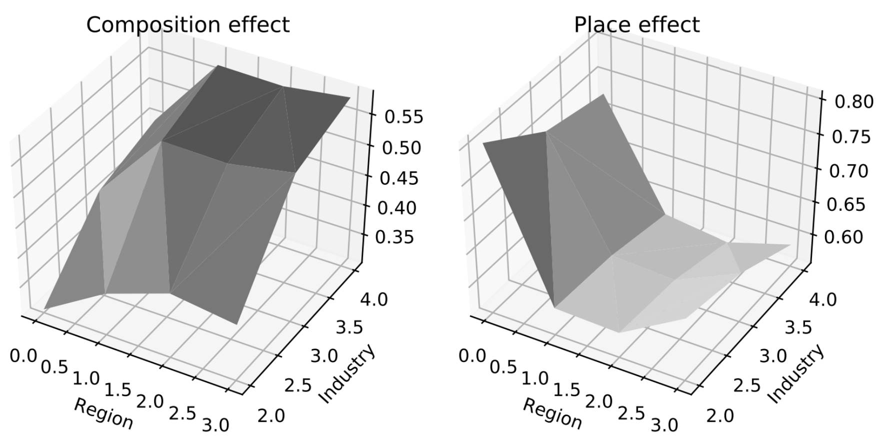

Composition (left) and Place (right) effects as functions of the level of aggregation in the industrial and geographical dimensions. The axis labeled as 'industry' measures the number of digits in NACE classification used for the decomposition routine. Equally, the axis labeled as 'region' indicates the number of digits in the NUTS classification, where NUTS0 is the national level.

I then attempt to take into account the subnational geographical heterogeneity in the proposed decomposition method. The Amadeus data used in this analysis contains an identifier for the 3-digit NUTS geographical location of the firm. We can, thus, estimate how the Composition and Place effects change when we change the geographical dimension from the most aggregate (NUTS0, which corresponds to nations) to the most disaggregate (NUTS3, which typically corresponds to cities and their broad surroundings). I find that disaggregation at the geographical level changes the estimates of both Composition and Place effects, with the former increasing and the latter dropping markedly. At the most disaggregated level I find that Composition and Place are equally important: 0.581 and 0.589 respectively. This prompts us to consider what is the effect of aggregation for the industrial dimension as well. The concern is even larger than in the geographical dimension. One in fact could argue that, although ignoring the within-nation geographical differences clearly hides some of the dynamics of productivity, a researcher could be interested to learn about decomposition at the country level. This is not the case for the industrial dimension where aggregation only gives biased information about productivity. In appendix C, a simple model of aggregation is proposed to show that industry disaggregation indeed corrects for noise, and therefore, disaggregated data are preferable.[12] Because of the conclusions of this model, in the previous section I only presented results at the maximum disaggregation possible in this data (4-digit NACE). I however test here with more aggregate industrial classification and find, in line with the aggregation model, that the importance of Composition drop, while that of the Place effect rises. The effect of aggregation on the estimates is summarized in figure 4. It is clear from the figure that disaggregating in both industry and regional dimensions increases, albeit non-monotonically, the importance of the Composition effect and decreases that of Place. Disaggregation – the figure indicates – changes

[12] The model also shows that disaggregation makes the decomposition terms converge to a finite limit and therefore cannot be used to obtain arbitrary values for the Composition and Place effects.

the conclusions deeply. Whereas in the aggregated case one would conclude that we live in a world dominated by the Place effect, Composition and Place are equally important if the regional and sectoral granularity is taken into account.

5.3 Technology diffusion

While disaggregation of data reveals concealed aspects of productivity, a case can be made that slicing the data along the industrial and regional dimension can provide new information as well. If, for instance, we compute the decomposition routine for east and west Europe separately, I argue that the difference can be interpreted as an indication of technological homogeneity within the group of countries. In fact, according to this interpretation, a stronger Composition effect signals greater homogeneity in industry technology across places, that is $\{A_{ig} - A_i\}$ is small for many industries. By this logic, one can interpret the difference in the Composition effect as an indicator of how easily technology diffuses. The argument is valid for both the geographical and industrial aggregation of data.[13]

Figure 5 shows some selected decomposition on subsets of the data. The decomposition of European countries is very similar to the benchmark, perhaps because the majority of the sample consists of countries in Europe. However, it can be observed that in the European Union, the Composition effect is approximately 50% larger than in the benchmark case, while in western Europe, it is almost twice as large, indicating greater technological homogeneity among these nations. In the subsample of Asian countries, the Composition and Place effects are almost equally important, although this must be interpreted with caution, as we have a very small subset of Asian countries (Japan,

[13] Though the industry dimension is not symmetric to the geographical one: the decomposition is a way to assess homogeneity of industries across geographies. By disaggregating, say, agriculture and manufacturing into several products we assess the homogeneity of productivity of potatoes and textile across geographies (rather than the homogeneity of technology between potatoes and textile).

South Korea, Australia, India, Hong Kong).

Figure 5: Decomposition of selected geographical or industrial subsets.

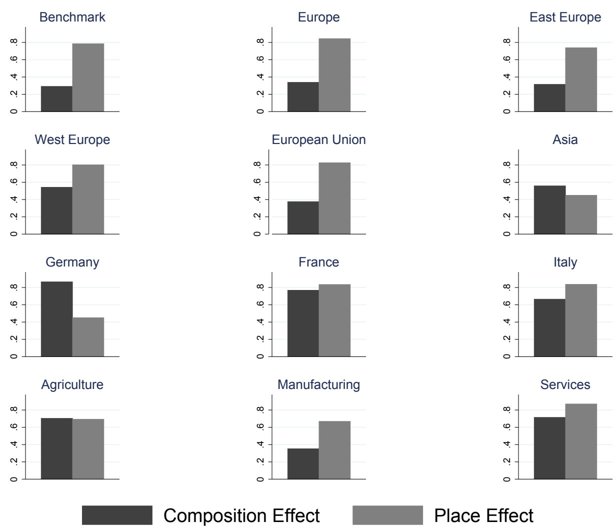

The standard assumptions on TFPQ used in the benchmark case apply. The Composition and Place effects are measured using the main indicator of relative standard deviation

If we look subnationally, at NUTS2 regions within a country, I find that on average the effect of Composition and that of Place are very similar, as shown in the productivity decomposition of France, Germany and Italy in figure 5.[14] While perhaps part of the larger impact of the Composition effect is due to lower level of noise in more disaggregated data, it is likely that some of the difference arises from less constrained knowledge diffusion *within* a country compared to *between* countries. Likely, technological homogeneity is easier to achieve at shorter distance and with a homogeneous population. To informally test the conjecture that the relative importance of the Composition and Place effects can be interpreted as a sign of the technological homogeneity, the following econometric specification is tested. The 16 European countries for which sufficient data are available to to run the decomposition analysis at

[14]With the German case exhibiting a Composition effect far stronger than the Place effect.

the subnational (NUTS2) level are selected.[15]

The decomposition analysis is run on country pairs; for instance, France and Italy would have 22 and 21 NUTS2 regions, respectively. Thus, the decomposition analysis for this pair of countries is limited to 43 NUTS2 regions. There are 120 country pairs ($(16^2-16)/2 = 120$), and for each of them, the ratio between the standard deviation of the Composition and Place effects, σ_C/σ_P, is computed. Table 4 reports the econometric results when this indicator is regressed against distance and contiguity, variables that refer to the 120 pairs of countries.

As one should expect, distance is negatively related to the ratio σ_C/σ_P, as countries that are farther apart have more heterogeneous technology. The effect, however, is significant only if a contiguity variable is not included. This could be interpreted as evidence that a short distance is not sufficient for knowledge diffusion, but the right channel is necessary. Possibly, contiguity allows for human capital migration, while distance without contiguity only allows for trade. Yet another interpretation is that a shared border indicates institutional similarity, which in turn facilitates technological diffusion. This would explain why sharing a border is more relevant than proximity and why the Composition effect is stronger between regions than between countries. The evidence presented, however, only allows for speculation, which is outside the scope of this analysis. The main point highlighted by these regressions is that the interpretation proposed – the variation in the relative importance of the Composition and Place effects can be seen as an indicator of technological homogeneity – is in line with the evidence on how σ_C/σ_P varies with distance and contiguity.

[15]These countries are Belgium, Bulgaria, the Czech Republic, Germany, Denmark, Spain, Finland, France, Hungary, Ireland, Italy, Poland, Portugal, Romania, Sweden, Slovakia.

Table 4: Distance decay of technology diffusion.

	Dependent variable: σ_C/σ_P			
	(1)	(2)	(3)	(4)
Distance	-0.087*		-0.040	-0.042
	(0.051)		(0.058)	(0.058)
Contiguity		0.230**	0.191*	0.202*
		(0.100)	(0.115)	(0.115)
GDP per capita (origin)				0.002
				(0.002)
GDP per capita (destination)				0.003
				(0.002)
Constant	1.032***	0.880***	0.941***	0.812***
	(0.079)	(0.037)	(0.095)	(0.127)
$Adj.R^2$	0.016	0.035	0.030	0.034
Obs.	120	120	120	120

Each observation represents a country pair. The dependent variable is obtained from the decomposition at the NUTS2 level of all regions in the pair. Standard errors are in parentheses, and significance is denoted as follows: $P < 0.1$*:*,*$P < 0.05$*:**,* $P < 0.01$*:****

A final comment, possibly hinting at future research, is on the subsets of the industrial dimension. Although the priors were less strong than for the other subsets shown in figure 5, it was somewhat unexpected that manufacturing had a smaller Composition effect than agriculture and services did. The evidence in Caselli (2005) and Rodrik (2013), in fact, suggests that manufacturing has a more homogeneous technology than agriculture (the former) or services (the latter).[16] The sectoral heterogeneity in technology diffusion is largely an understudied topic. Yet, its importance is far from marginal because – as highlighted in section 1 – the existence of asymmetric externalities between industries has severe implications for which type of world we live in.

[16]Rodrik (2013) argues that technological convergence is easier in manufacturing because it is more tradable and, hence, subject to stronger competition.

For this reason, the contradictory evidence cannot lead to swift conclusions but only calls for further research.

6 Conclusions

The estimates in this paper suggest that the difference in productivity across a sample of emerging and developed economies is due to both technological differences (Place effect) and allocation of resources (Composition effect). A conservative approach suggests that the former is responsible for approximately 1/4 of the differences and the latter for 3/4. Allowing for geographical differences within-country strengthen this results even further, with Place and Composition each explaining around 1/2 of the differences in productivity.

These estimates imply that both reducing the technological gap and reallocating resources have the potential to improve lagging countries' productivity and hence reduce the cross-national differences in standards of living. Technology adoption, given its central role, appears to be a *conditio sine qua non* for economic development. Nevertheless, it might prove insufficient to achieve full catch up, as without reallocation of labor and capital into more productive activities, even full technological catch-up (and taking the conservative estimates) would still leave 29% of the difference in productivity. This hints that the role of structural change might be far from marginal in a country's growth process.

From an academic perspective, this result might add to the recent stream of empirical and theoretical articles (Imbs and Wacziarg, 2003; Hausmann et al., 2007; Cadot et al., 2011) showing that there is room for a revised role of structural change in the theory of growth, a role that used to be prominent in earlier theories of development. From a policy perspective, the analysis suggests lagging countries need to narrow a productivity gap, while also solving an allocation problem. The reason for this problem of composition might at least in part be technological and, hence, mitigated once the productivity gap is narrowed. However, institutional, infrastructural and demand-side factors can create a situation in which a market-winning technology is available, but a country fails to specialize in it. In Diodato, Malerba, and Morrison (2018a), for instance, a model shows how demand externalities might lead to failures in sectoral catch-up, even after the lagging country has entirely closed the productivity gap. This suggests that the adoption of an industry-specific technology does not directly imply specialization in that industry. Future research on the factors that govern the link between adoption and diversification, as well as on factors that influence structural change – intended as a reallocation of factors of production from low productivity to high productivity industries – might improve our still imperfect understanding of economic growth.

References

Banerjee, Abhijit V. and Esther Duflo (2005), "Growth theory through the lens of development economics." In *Handbook of Economic Growth 1A* (Philippe Aghion and Steven N. Durlauf, eds.), 473–552, Elsevier.

Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta (2013), "Cross-country differences in productivity: The role of allocation and selection." *The American Economic Review*, 103, 305–334.

Cadot, Olivier, Céline Carrere, and Vanessa Strauss-Kahn (2011), "Export diversification: what's behind the hump?" *The Review of Economics and Statistics*, 93, 590–605.

Caselli, Francesco (2005), "Accounting for cross-country income differences." In *Handbook of economic growth 1A* (Philippe Aghion and Steven N. Durlauf, eds.), 679–741, Elsevier.

Diodato, Dario, Franco Malerba, and Andrea Morrison (2018a), "The made-in effect and leapfrogging: A model of leadership change for products with country-of-origin bias." *European Economic Review*, 101, 297–329.

Diodato, Dario, Frank Neffke, and Neave O'Clery (2018b), "Why do industries coagglomerate? How Marshallian externalities differ by industry and have evolved over time." *Journal of Urban Economics*, 106, 1–26.

Ellison, Glenn, Edward L. Glaeser, and William R. Kerr (2010), "What causes industry agglomeration? Evidence from coagglomeration patterns." *The American Economic Review*, 100, 1195–1213.

Galor, Oded and Andrew Mountford (2006), "Trade and the great divergence: The family connection." *The American Economic Review*, 96, 299–303.

Hausmann, Ricardo, Cesar Hidalgo, Daniel P Stock, and Muhammed Ali Yildirim (2014), "Implied comparative advantage." *HKS Working Paper No. RWP14-003.*

Hausmann, Ricardo, Jason Hwang, and Dani Rodrik (2007), "What you export matters." *Journal of Economic Growth*, 1, 1–25.

Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi (2014), "Growth and structural transformation." In *Handbook of Economic Growth 2B* (Philippe Aghion and Steven N. Durlauf, eds.), 855–941, Elsevier.

Hsieh, Chang-Tai and Peter J Klenow (2009), "Misallocation and manufacturing TFP in China and India." *The Quarterly Journal of Economics*, 124, 1403–1448.

Hsieh, Chang-Tai and Peter J Klenow (2010), "Development accounting." *American Economic Journal: Macroeconomics*, 2, 207–223.

Imbs, Jean and Romain Wacziarg (2003), "Stages of diversification." *American Economic Review*, 63–86.

Kuznets, Simon (1966), *Modern economic growth*. Yale University Press New Haven.

Levinsohn, James and Amil Petrin (2003), "Estimating production functions using inputs to control for unobservables." *The Review of Economic Studies*, 70, 317–341.

Matsuyama, Kiminori (1992), "Agricultural productivity, comparative advantage, and economic growth." *Journal of Economic Theory*, 58, 317–334.

McMillan, Margaret, Dani Rodrik, and Íñigo Verduzco-Gallo (2014), "Globalization, structural change, and productivity growth, with an update on Africa." *World Development*, 63, 11–32.

Olley, G Steven and Ariel Pakes (1996), "The dynamics of productivity in the telecommunications equipment industry." *Econometrica*, 64, 1263–1297.

Porter, Michael (2003), "The economic performance of regions." *Regional Studies*, 37, 545–546.

Rodrik, Dani (2013), "Unconditional convergence in manufacturing." *The Quarterly Journal of Economics*, 128, 165–204.

Timmer, Marcel P and Gaaitzen J de Vries (2009), "Structural change and growth accelerations in Asia and Latin America: a new sectoral data set." *Cliometrica*, 3, 165–190.

Young, Alwyn (1993), "Invention and bounded learning by doing." *Journal of Political Economy*, 101, 443–472.

Appendix

A Elasticity of substitution

To move from TFPR to TFPQ, this paper uses the HK correction, which is obtained by elevating value added to the power of $\sigma/(\sigma-1)$. Hsieh and Klenow (2009) set $\sigma = 3$, as they note that the gains from liberalization – the objective of their analysis – are increasing in σ and setting it low results in a conservative estimate. If we were to apply the same reasoning, then we would have to set a high σ: our prior is a low Composition effect, and a conservative estimate would need to penalize it. Although the sensitivity to σ is not monotonic (see figure A.2), on average, the Composition effect is decreasing in σ. In spite of this argument, We do not find it desirable to set an excessively high value for the benchmark case. As $\sigma/(\sigma-1)$ approaches unity, the proposed measure of productivity approaches TFPR, and it has been stressed – in the literature and in section 3 of this paper – that the correct measure is TFPQ. The benchmark value ($\sigma = 7$) is chosen with the aid of figure A.1.

Figure A.1 shows that changing the value of σ significantly affects the ranking of industries by productivity. However, while above a value of 6, the differences are contained, below this value, the rankings become highly unstable – to the extent that for $\sigma = 3$, one of the bottom-10 industries becomes more productive than *telecommunications* – the most productive for a larger elasticity.

Figure A.1: Changes in ranking of Top-10 and Bottom-10 industries by TFPQ.

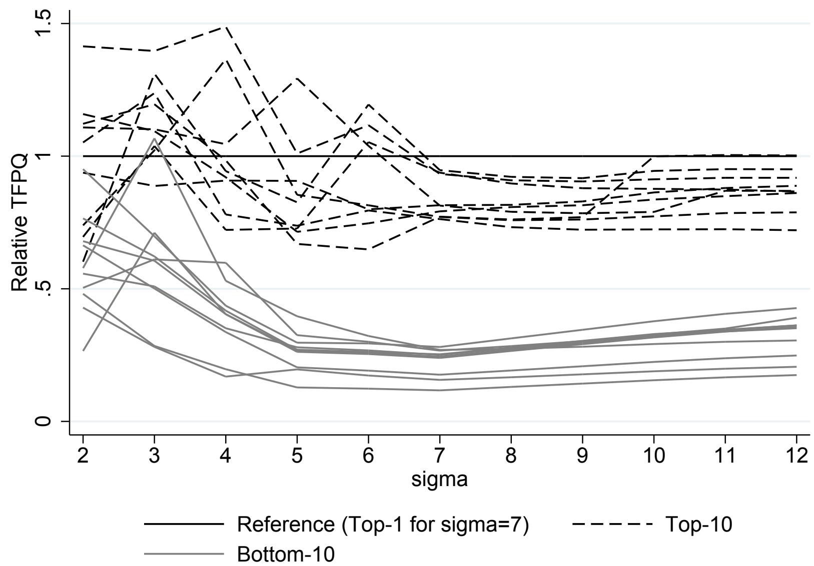

All productivities are expressed as TFPQ, relative to the most productive industry for $\sigma = 7$ *(telecommunications).*

If, on the one hand, a high σ is undesirable because it is closer to TFPR than TFPQ, on the other hand, a low σ is also also undesirable because of ranking instability. The value of 7 is chosen not only because it is a compromise between a low and a high elasticity, but also because it is the closest value to the reference value in the literature (Hsieh and Klenow, 2009) that maintains ranking stability. Additionally, as shown in figure A.2, departing from $\sigma = 7$ does not create major changes in the Composition and Place effects, which remain relatively stable around the benchmark values.

Figure A.2: Sensitivity of Composition and Place effects to the elasticity of substitution.

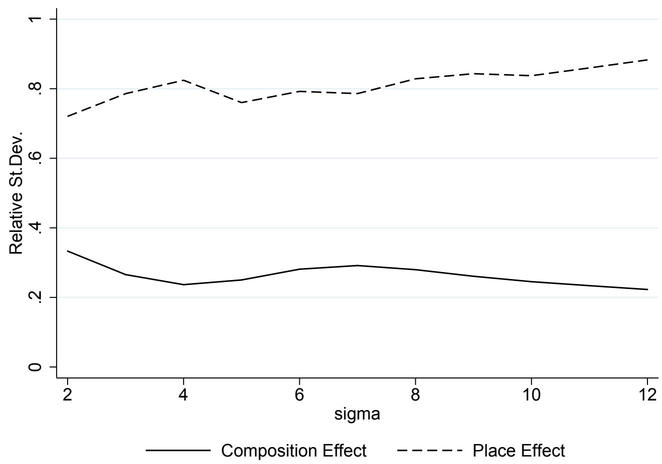

Composition and Place effects are measured with the standard indicator.

B Share of capital and labor

This section discusses the sensitivity of the results to the estimation method. It is shown that, although FE and LP estimation increase the estimates of both the Composition and Place effects, the results are qualitatively similar. Table B.1 provides a brief summary. Since for these estimates only a restricted sample of European countries is used, the first column of table B.1 – reporting the standard case, with $\alpha = 0.33$ and $\beta = 0.67$ – shows marginally larger values for the Composition effect and marginally lower values for the Place effect, compared to the benchmark case in table 3.

Table B.1: Estimates of Capital and Labor coefficients.

	(Standard)	(OLS)	(FE)	(LP)
Capital share (α)	0.33	0.324***	0.182***	0.137***
		(0.000)	(0.001)	(0.003)
Labor share (β)	0.67	0.608***	0.588***	0.522***
		(0.001)	(0.001)	(0.002)
R^2		0.659	0.249	.
Obs.		1345691	1345691	1290754
Composition	0.367	0.387	0.541	0.728
Place	0.769	0.796	0.885	0.932

For the first column, α and β are assumed. The upper part of the table reports the econometric estimates. For FE, the R^2 row reports within R^2. Standard errors are in parentheses, and significance is denoted as follows: $P < 0.1$:,$P < 0.05$:**, $P < 0.01$:***. The lower part of the table displays the corresponding Composition and Place effects.*

The OLS estimates reported in the second column are a strong justification for the choice of parameter values in the standard case. In fact, the estimated value of $alpha$ is almost exactly 1/3, while that of β is only a few percentage points below the assumed value of 2/3. Subsequently, the decomposition results are very close to the standard case. For alternative estimation techniques, results differ more from the standard case. The exponent of labor is not changed in the fixed effects (FE) estimates, but the capital share is significantly reduced. This has the consequence of increasing *both* the Composition effect and Place effect by 10 percentage points. However, since they both increase, a qualitative assessment of the importance of the two terms is not affected by these changes. More extreme is the deviation from the benchmark results under LP estimation. The estimated coefficients of capital and labor are much lower. The resulting decomposition suggests that both the Composition and Place effects are considerable more important, to the extent that the latter is approaching unity. The Composition effect is also markedly more important in this case, with the indicator being almost twice as large as in the standard case. One interpretation of such large values is as follows: a reduction in the difference in composition ($s_{ig} = s_i$) would lead to a standard deviation of average productivity that is only 7 percentage points lower than before, hence only marginally reducing the productivity differences. Similarly, reducing differences in productivity ($A_{ig} = A_i$) would lead to a standard deviation of average productivity that is 70% of the previous value. This is quite larger than before and would imply a lower impact of equalizing industry productivity. In summary, the importance of both the Composition and Place effects changes drastically in such an estimation. Nonetheless, the qualitative conclusion that both effects are important, with the Place effect capturing the largest part of the variance, still holds.

C Aggregation model

This appendix presents a stylized model of industry aggregation. The aim of this model is twofold: first, to show that aggregation noise exists. Second, it demonstrates one cannot obtain arbitrary values of the Composition and Place effects by selecting a level of disaggregation of his own choosing: it is shown that under some (reasonable) assumptions, a 'true' value of the Composition effect exists.[17]

It is rather straightforward to show the worrisome outcome of a model in which the Composition effect ranges from 0 to 100% of the variance in productivity, according to how aggregated the data are: if the entire economy is considered to be a single industry, then there would be only a Place effect and no Composition effect. If each firm had a different industry code, then there would only be a Composition effect and no Place effect. This model, however,

[17] As highlighted in section 5.2, the argument that aggregation creates noise is stronger in the industry case. We only develop, therefore, a model of industry aggregation, although aggregation certainly also creates noise in the geographical dimension. Moreover, to further simplify the analysis, only the Composition effect is studied.

uses the implicit assumption that each firm is entirely unique. This assumption is absurd in the context of this research, as two firms would never have the possibility of learning from one another (no technology diffusion would be possible in this world). Moreover, and perhaps more important, if each firm had a different industry code, we would lose the ability to compare countries (each country would have its own set of industries). A more reasonable assumption in constructing a model of aggregation is to think of firms as belonging to a continuum of industries. In this way, it is possible to infinitely disaggregate to more fine-grained classifications, while still maintaining a logical link of similarity between firms. For a location g, let us imagine that labor and production are distributed across industries over a continuum $i \in [0, 1]$. We write the following convenient distributions.

$$L_g(i) = \frac{1}{Q-1} Q^i ln(Q) L_g, \quad Y_g(i) = \frac{1}{Z-1} Z^i ln(Z) Y_g, \tag{C.1}$$

where Q and Z are random variables (each location draws a value) with expected values q and z. Note that $\int_0^1 L_g(i)\, di = L_g$ and $\int_0^1 Y_g(i)\, di = Y_g$. To further clarify the meaning of these distributions, let us assume that a location draws the following values: $Q = 2$ and $Z = 3$. Figure C.1 depicts the following quantities $L_g(i)$, $Y_g(i)$ and $A_g(i) = Y_g(i)/L_g(i)$[18] for this example. Total labor and output in the country are normalized to 1.

The figure is intended to synthesize how input and output are distributed across the continuum of industries ranging from $i = 0$ to $i = 1$. Subsequently, this also provides information about the efficiency at which industry i transforms input into output. It is important to highlight that, in this example, as Q and Z are drawn with values larger than one, the industries with index i close to one have larger shares than those with index i close to zero.

[18] For simplicity, the stylized model of aggregation presented here considers labor productivity instead of TFP.

Figure C.1: Distribution of labor and output across a continuum of industries.

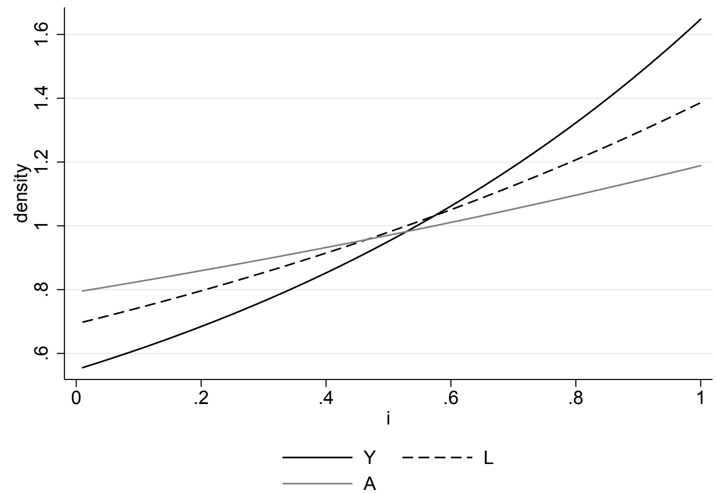

The figure depicts an example for one location. In this example, the location extracted $Q = 2$ and $Z = 3$. As $Z > Q$, $A := Y/L$ slopes positively, which means that, for this realization, more productive activities have larger employment shares.

Moreover, since $Z > Q$ (with $Q > 1$), the industries that have larger shares are also the most productive (i.e., there is a positive covariance between the size of an industry and its productivity). Conversely, if $Z < Q$, this means that the country is specialized in less-productive industries. By tweaking the assumptions on the probability distribution of Z and Q and the overall productivity parameter Y_g/L_g, one can recreate the different worlds we are attempting to analyze. For instance, with little covariance between Q and Z and systematic differences in the average productivity Y_g/L_g, one can recreate a scenario in which the Place effect dominates. Conversely, smaller differences in Y_g/L_g and large $cov(Z,Q)$ can be set to mimic a world in which the Composition effect is the most important contributor to the differences in productivity.

The point of this model is not to propose an effective way to describe the underlying process behind the differences in productivity. Its purpose is instead to supply a handy tool for understanding what happens when we use the indicators of this analysis on more aggregated data. In fact, the distributions in C.1 are chosen because they can be conveniently used to describe the aggregation of industries into larger groups. Let us imagine that we want to aggregate the employment of country g in I industries, starting from a continuum between zero and one. We are, in other words, transforming – discretizing – $L_g(i)$ into L_{ig}, where $i \in I$. This operation can be written as

$$L_{ig} = \int_{(i-1)/I}^{i/I} L_g(i)\, di = Q^{i/I}\left(\frac{1 - Q^{-1/I}}{Q - 1}\right) L_g. \tag{C.2}$$

The same operation can be carried out for production.

$$Y_{ig} = \int_{(i-1)/I}^{i/I} Y_g(i) di = Z^{i/I}\left(\frac{1 - Z^{-1/I}}{Z - 1}\right) Y_g. \tag{C.3}$$

It can be verified that the estimate of a country's productivity, A_g, is not influenced by aggregation. This is not the case, however, for the decomposition indicators used in this paper. As expected, the transformed indicators are a function of I – which specifies how coarsely the economy is aggregated into discrete industries. The Composition term, for instance, is[19]

$$C_g = \sum_{i=1}^{I} \frac{L_{ig}}{L_g} \frac{Y_i}{L_i}. \tag{C.4}$$

The Composition term is obtained by substitution

$$C_g = \sum_{i=1}^{I} \frac{(1 - Q^{-1/I})(1 - z^{-1/I})(q - 1)}{(1 - q^{-1/I})(Q - 1)(z - 1)} \left(\frac{Qz}{q}\right)^{i/I} \frac{Y}{L}. \tag{C.5}$$

In equation C.5, everything except the term $(Qz/q)^{i/I}$ is constant across in-

[19]It is assumed here that, as Q is uniformly distributed in a narrow range around its expected value, it is possible to treat the distribution of L_{ig} as linear and approximate L_i as $L_i = E[L_{ig}(Q, L_g)]G = L_{ig}(E[Q], E(L_g))G$, where G is number of locations. That is, $L_i = q^i ln(q)L/(q - 1)$. Similarly, we can write $Y_i = z^i ln(z)Y/(z - 1)$. Numerical simulations (available upon request) show that this approximation only creates small errors ($< 1\%$ for $q = 2$, $< 0.001\%$ for $q = 100$).

dustries (i) and can be taken out of the sum. The remaining term, $\sum_{i=1}^{I}(Qz/q)^{i/I}$, is a geometric series and is equal to

$$\sum_{i=1}^{I}\left(\frac{Qz}{q}\right)^{i/I} = \frac{\left(\frac{Qz}{q}\right)^{i/I}\left(\frac{Qz}{q}-1\right)}{\left(\frac{Qz}{q}\right)^{i/I}-1}. \tag{C.6}$$

With minor algebra, it is possible to find the following expression for the Composition term, C_g, as a function (ω) of I, Q, z and q

$$C_g := \omega(I,Q,z,q) = \underbrace{\frac{Y}{L}\frac{(q-1)\left((Qz/q)-1\right)}{(Q-1)(z-1)}}_{\psi}\underbrace{\frac{(Q^{1/I}-1)(z^{1/I}-1)}{(q^{1/I}-1)\left((Qz/q)^{1/I}-1\right)}}_{\lambda(I)}. \tag{C.7}$$

Our interest lies in understanding what happens if $I \to \infty$. Since ψ, the first part of equation C.7, does not depend on I, we can treat it as a constant. To solve the limit, it is sufficient to split $\lambda(I)$, the second part of equation C.7, into two terms [a] $(Q^{1/I}-1)/(q^{1/I}-1)$ and [b] $(z^{1/I}-1)/((Qz/q)^{1/I}-1)$. Each is solved using L'Hôpital's rule. The limit of ω is then equal to

$$\lim_{I\to\infty}\omega(I,Q,z,q) = \psi\frac{ln(Q)ln(z)}{ln(q)ln(Qz/q)}. \tag{C.8}$$

The fact that ω has a finite limit for $I \to \infty$ addresses my concern: the outcome of the decomposition does not arbitrarily change with changes in the industry aggregation. Instead, it changes by approaching its 'true' value when data are sufficiently disaggregated.