Oxford Bulletin of Economics and Statistics

ORIGINAL ARTICLE OPEN ACCESS

Regional Productivity Differences in the UK and France: From the Micro to the Macro

Bridget Kauma[1] | Giordano Mion[2,3,4]

[1]Department of Economics, Business School, University of Sussex, Brighton, UK | [2]Department of Economics, ESSEC Business School, Cergy, France | [3]School of Economics, University of Nottingham, Nottingham, UK | [4]Centre for Economic Performance, The Centre for Economic Policy Research, the CESifo Research Network, Munich, Germany, The National Institute of Economic and Social Research and THEMA, London, UK

Correspondence: Giordano Mion (giordano.mion@essec.edu)

Received: 31 October 2023 | **Revised:** 15 May 2025 | **Accepted:** 20 May 2025

Funding: This work was supported by Labex MME-DII and Productivity Institute.

Keywords: density | firm-level dataset | France | Markups | productivity | regional disparities | UK

ABSTRACT

We propose a new data resource that attempts to overcome limitations of standard firm-level datasets for the United Kingdom (like the ARD/ABS) by building on administrative data covering the population of UK firms with at least one employee. We also construct a similar dataset for France and use both datasets to (1) provide some highlights of the data and an overall picture of the evolution of aggregate UK and French productivity and markups; (2) analyse the spatial distribution of productivity in both countries at a fine level of detail—228 Travel to Work Areas (TTWAs) for the United Kingdom and 297 Zones d'emploi (ZEs) for France—while focusing on the role of economic density. Our findings suggest that differences in firm productivity across regions are magnified in the aggregate by an increasing productivity return of density along the productivity distribution.

JEL Classification: R12, D24

1 | Introduction

A stylised fact of economic geography is that the productivity of firms increases with city size and urban density [1]. A large literature dating back to Marshall [2] explores the question of why cities have this productivity advantage. Micro-foundations put forward for these agglomeration externalities are typically grouped under the headings of sharing, matching, learning and sorting [3, 4] and include different forms of knowledge spillovers between firms, costly trade, pro-competitive effects of city size and sorting of workers [5]. The empirical literature suggests a rather consistent, across countries and years, range for the elasticity of productivity with respect to city size. Rosenthal and Strange [6] and Combes and Gobillon [1] provide summaries of this literature and agree on a range for the key elasticity of productivity with respect to density of 0.02–0.10.[1]

While most geographers would typically consider regions as the unit of analysis and directly work at this level of aggregation, economists are increasingly using firms or even establishments as the unit of analysis around which to reconstruct and attribute differences in economic performance across regions. Crucially, the two approaches do not seem to provide the same magnitudes regarding, for example the elasticity of productivity with respect to local density. More specifically, Jacob and Mion [9] provide evidence for French manufacturing firms highlighting the importance of weighting in going from the firm-level (micro) to the regional-level (macro) productivity. They find smaller values for the elasticity of productivity with respect to population density when using unweighted firm-level regressions while getting somewhat larger values when considering revenue- or employment-weighted firm-level regressions.

The productivity of a region clearly reflects the productivity of its firms. However, the aggregate productivity of a region is the weighted average (typically by employment) of the productivity of the firms located in the region and not the simple average. When running unweighted regressions, in which firms are the unit of analysis, to measure productivity differences across space one is essentially comparing the average firm across different locations irrespective of the firm size distribution and its link to productivity within regions. The link between micro and macro is restored if one runs weighted regressions, and the coefficients from the unweighted and weighted regressions do not need to be the same. One reason they could differ is a varying (across regions) correlation between firm size and firm productivity. For example, if denser regions are characterised by a higher correlation between firm size and firm productivity, unweighted differences in productivity across space will be magnified when weighting. Another reason for differences between coefficients is the heterogeneity (along the productivity dimension) of the elasticity of productivity with respect to density. For example, if more productive firms enjoy disproportionate gains from the density of economic activities, unweighted differences in productivity across space will again be magnified when weighting because (on average) more productive firms are larger.

In what follows, we extend the analysis of Jacob and Mion [9] beyond manufacturing to the whole private sector for both France and the United Kingdom, while digging into the above-mentioned explanations for the larger values of the elasticity of productivity with respect to population density when using weighted firm-level regressions as compared to unweighted firm-level regressions. In order to achieve this, we first construct two large datasets spanning the entire population of French and UK firms with at least one employee, allowing us to retrieve different measures of productivity – including labour productivity and total factor productivity (TFP) – and investigate the links between productivity and geography at a fine spatial level: 228 Travel to Work Areas (TTWAs) for the UK and 297 Zones d'emploi (ZEs) for France. Considering the last year of the data, 2017, the datasets we constructed span over 814,407 firms employing 17,441,714 workers for the United Kingdom and over 900,026 firms employing 12,406,277 workers for France. In both cases, the availability of the location of the different establishments belonging to each firm allows us to link productivity to space and perform our investigations.

While the French data we build upon here has been used in many other studies in the past, the longitudinal dataset we have constructed for the United Kingdom has never been assembled before and this is one of our key contributions.[2] The ARD/ABS surveys administrated by the ONS have so far been used to study, for example firm productivity in the United Kingdom. However, the key advantage of our data is the much higher coverage of firms allowing us, for example to study the spatial distribution of productivity at a fine spatial level like TTWAs.

After describing the construction of the two datasets, we provide some data highlights regarding productivity, markups and the financial crisis period, which is included in the time span of our analysis (2000–2017 for France and 2004–2017 for the United Kingdom). We provide various comparable metrics, based on a similar underlying exhaustive data source, on the level and evolution of productivity and markups for both France and the United Kingdom. This is another contribution of our paper. Considering the United Kingdom, while total factor productivity has been both only very lightly and very briefly affected by the financial crisis, the same is not true for markups, revenue per worker and labour productivity, which is consistent with evidence provided in analyses based on the smaller ARD/ABS datasets like Harris and Moffat [11] and Jacob and Mion [12]. Inspection of markups reveals that they recovered their pre-financial crisis level around 2015, while for labour productivity, the recovery year is 2016. As for France, it is not entirely clear whether total factor productivity had by 2017 picked up its pre-financial crisis level. On the other hand, revenue per worker and labour productivity have been little affected by the financial crisis. Inspection of markups reveals that they have not yet recovered their pre-financial crisis level, suggesting that French firms struggle to achieve pre-financial crisis profit margins.

Turning to the spatial analysis (our third contribution), in our investigation, we primarily focus on 'single region firms', those firms that we can uniquely associate with one region. Such firms may thus have more than one establishment, but such establishments need to be located in the same region. The reason we are particularly interested in single-region firms is that for such firms, there are no issues in, for example attributing their productivity and their employment to a particular region. Single-region firms represent the vast majority of firms and account for about half of the overall employment. We also provide some robustness results, including multi-region firms in the analysis, while attributing the same productivity to all of the establishments of a given multi-region firm and using establishment-level employment for weighting. Such robustness results largely uphold our findings based on the single-region firms.

Our results can be summarised as follows: First, for both France and the United Kingdom, we find a larger productivity return to density when weighting observations by employment as compared to unweighted regressions. Digging deeper into this reveals, in both cases, the following: (1) The correlation between firm size and productivity within a region is quite low (and sometimes negative) across regions, particularly for the United Kingdom. (2) The relationship between these correlations and regional density is not positive and actually slightly negative. These findings indicate that, if anything, the varying correlation between firm size and productivity across regions should *reduce* and not amplify spatial productivity differences when going from the micro to the macro. On the other hand, in both cases, we find evidence that the productivity return of density is increasing along the deciles of the productivity distribution. This finding is reminiscent of the 'dilating' of the productivity distribution in larger regions found for France in Combes et al. [13]. It is such an increasing productivity return of density that magnifies firm-level productivity differences for the United Kingdom and France, when going from the micro to the macro, and not a varying correlation between firm size and productivity across space.

In terms of the comparison between the United Kingdom and France, we find the following: In the United Kingdom, the correlation between firm size and productivity within a region is low (compared to France) and sometimes negative. If the United Kingdom had French correlations, aggregate productivity would be higher. Also, the United Kingdom has a problem of productivity being quite unequal across space beyond density (big London gap, while little Paris gap). The problem with France is instead the negative productivity return to density for the least productive firms, that is denser places in France nurture too many low productive firms, and this creates a big divide between the unweighted and weighted productivity return to density.

The rest of the paper is organised as follows: Sections 2 and 3 present the data sources we use and describe how we cleaned and combined the data together for the United Kingdom and France, respectively. Section 4 provides details of the productivity and markups estimations, while Section 5 presents some data highlights and an overall picture of the evolution of aggregate UK and French productivity and markups. Section 6 delineates our conceptual framework, while Section 7 contains our spatial analysis. Section 8 provides instead a number of robustness exercises. Section 9 concludes. Additional details about the data are provided in Appendix A while complementary tables and figures are reported in the Supporting Information.

2 | UK Data

2.1 | Data Sources

2.1.1 | BSD

The Business Structure Database (BSD) is an annual extract (the snapshot taking place at the end of a fiscal year) of the Inter-department Business Register (IDBR), a live database of business organisations in the United Kingdom. Organisations that are registered for VAT or pay at least one member of staff through the Pay As You Earn (PAYE) tax system will appear on this register.

The BSD is administered by the ONS and, while being one of the largest sources of data about business organisations in the United Kingdom, it contains only a limited number of variables. In our analysis, we borrow information about the number of employees, employment (number of employees plus owner(s)) and foreign ownership. A firm in the BSD is identified by a unique code, which we refer to as the 'BSD firm id'. The BSD also provides information on the employment and location (up to the postcode level) of the different establishments belonging to a given firm that we use for our spatial analysis. An establishment in the BSD is identified by a unique code, which we refer to as the 'BSD establishment id'.

2.1.2 | Vat

The Value Added Tax (VAT) panel database is an annual extract from VAT returns providing information on organisations that are registered for VAT.

The VAT panel database is administered by HMRC and provides information on, among other things, the value of purchases operated in a given (fiscal) year as well as the value of sales. A firm in the VAT panel database is identified by her unique VAT code, which is anonymised within the HMRC datalab environment, and we refer to it as the 'VAT firm id'.

2.1.3 | Fame

FAME contains information on companies registered at Companies House in the United Kingdom. It covers company financials, corporate structures, shareholders and subsidiaries. The data are collected from various sources, most notably the national official bodies in charge of collecting company accounts data, and are then compiled and organised by Bureau van Dijk (BvD). FAME is available within the HMRC Datalab, where original company identifiers are anonymised.

The coverage of variables like sales, intermediate purchases and employment in FAME is very patchy because only relatively large firms are required to report this information in their annual accounts. However, information on assets, and in particular on tangible fixed assets, which we are going to use as our measure of the firm capital stock, is very well recorded. A firm in FAME is identified by its unique anonymised CHR number, which we refer to as the 'FAME firm id'.

2.2 | Cleaning and Combining the Data

In what follows, we explain how we cleaned and merged the data while relegating some details to Appendix A. The data are organised by fiscal year; for instance, the year 2017 refers to the fiscal year 2017–2018.

2.2.1 | Data Cleaning

BSD. For the BSD, we first worked on the industry classification to derive consistent information on the SIC 2007 primary code of each firm.[3] We have subsequently eliminated firms involved in financial and insurance activities (SIC 2007 codes 64, 65 and 66) and restricted the sample to firms with at least one employee and with a live VAT status.[4] A firm in the data is identified by the BSD firm id, and the data span from 2004 to 2017.

VAT. Again we applied some cleaning to the industry classification (which is time varying in the VAT panel dataset). Firms involved in financial and insurance activities are dropped from the analysis.[5] A firm in the data is identified by the VAT firm id and the data span from 2004 to 2017.

FAME. We cleaned the data from some duplicates and kept only observations for which the variable fixed assets are not missing.[6] A firm in the data is identified by the FAME firm id and the data span from 2004 to 2017.

2.2.2 | Data Matching

Each of the three datasets has a different firm identifier, and the correspondence between any pair of identifiers is, in some cases, many-to-many. The HMRC datalab provides a lookup table across the 3 identifiers, but the many-to-many correspondence issue still needs to be addressed. We explain in Appendix A how we solve this problem and end up using an 'aggregate' definition of a firm identified by what we label 'final firm id'.

2.2.3 | Adding Information on Location

In order to retrieve the location(s) of a firm we use the information on local units from the establishments files of the BSD. Each BSD establishment id is uniquely linked to a BSD firm id and so to a unique final firm id. For each final firm id in our data, we can then identify the related establishments and for each such establishment the BSD provides information on location (up to the postcode level) and employment. To allow a meaningful spatial analysis, we use an 'economic' partition of the UK geography and in particular the 2011 version of the Travel To Work Areas (TTWAs). The 2011 version of TTWAs breaks down the United Kingdom (including Northern Ireland) into 228 areas. See Appendix A for further details.

Equipped with the information above, we are thus able to identify what we refer to as 'single TTWA firms', those firms that we can uniquely associate with one TTWA. Such firms may thus have more than one establishment, but such establishments need to be located in the same TTWA. The reason we are particularly interested in single TTWA firms is that for such firms, there are no issues in, for example attributing their productivity and their employment to a particular TTWA. By contrast, for multi-TTWA firms (e.g., Tesco), it is less clear how to allocate productivity (which can be measured only at the level of the firm) to the different TTWAs in which the firm has establishments. Single TTWA firms represent the vast majority of firms (around 97%) and account for about 43% of overall employment in our dataset.

3 | French Data

3.1 | Data Sources

3.1.1 | FICUS

FICUS is an administrative database containing detailed accounting information (employment, sales, intermediates, capital, industry affiliation etc.) for the population of French firms. The database is part of the SUSE (Système unifié de statistiques d'entreprises) framework. SUSE constitutes a coherent set of statistical data on firms obtained from the joint use of two sources of information: tax declarations of companies to the General Directorate of Taxes (DGT) and the annual business surveys (EAE).

FICUS includes both balance sheet and profits and losses account information, and, for the purpose of our analysis, we use information from the year 2000 till the year 2007, when the dataset was replaced by the companion database FARE (see below). Each firm in the dataset is uniquely identified by a 9-digit code (SIREN code).

3.1.2 | FARE

FARE is an administrative database containing detailed accounting information (employment, sales, intermediates, capital, industry affiliation, etc.) for the population of French firms. The database is part of the ESANE (Élaboration des statistiques annuelles d'entreprises) framework. The ESANE framework succeeded the previous framework (SUSE) and, since 2008, this new system has jointly exploited, via a specific estimation procedure, administrative data and data from the ESA and EAP surveys in order to produce the most accurate sectoral statistics possible.

FARE includes both balance sheet and profits and losses account information and, for the purpose of our analysis, we use information from the year 2008 to the year 2017. Each firm in the dataset is uniquely identified by a 9-digit code (SIREN code).

3.1.3 | Stocks D'établissements

The Stocks d'Établissements database is a demography product of establishments providing identity data on the characteristics of establishments. It includes establishments active on 31 December of year N. The data are compiled from the Directory of Companies and Establishments (REE).

The Stocks d'Établissements database contains information on the location of establishments (up to the municipality level) as well as on their employment. Each establishment in the dataset is uniquely identified by a 14-digit code (SIRET code), which can be uniquely attached to a SIREN (firm) code.

3.2 | Cleaning and Combining the Data

In what follows, we explain how we cleaned and merged the data.

3.2.1 | Data Cleaning

FICUS and FARE. See Appendix A for details.

Stocks d'Établissements. Considering the Stocks d'Établissements dataset, we simply discard observations with missing SIRET and/or municipality code.

3.2.2 | Data Matching

Data matching is quite straightforward with French firm data because of the unique firm identifier (SIREN code). We thus simply append the information for each year, coming from either FICUS or FARE, thus obtaining a panel of firms over the period 2000–2017. Finally, we define industries as two-digit NACE rev2 codes and apply some grouping (detailed below) in preparation for TFP estimations.

3.2.3 | Adding Information on Location

Each establishment in the Stocks d'Établissements database is identified by a unique 14-digit code (SIRET code) whose first 9 digits correspond to the SIREN code of the firm. This greatly facilitates the task of adding location information.

To allow a meaningful spatial analysis we use, as in the UK case, an 'economic' partition of French geography and in particular the 2010 version of the Zone d'Emplois (ZEs), providing us with 297 areas for continental France.[7]

Equipped with the information above, we are thus able to identify, as in the case of the United Kingdom, single ZE firms, those firms that we can uniquely associate with one ZE. Single ZE firms represent the vast majority of firms (around 93%) and account for about 53% of overall employment in our dataset.

4 | Productivity and Markup Estimation

In order to estimate productivity and markups, we use a production function approach. For the United Kingdom, we use sales from the VAT data as a measure of output/revenue, purchases from the VAT data as a measure of intermediate expenditure, tangible fixed assets from FAME as a measure of the capital stock, and employment (count of employees plus the owner(s)) from the BSD as a measure of the labour input. For France, we use firm turnover as a measure of output/revenue, purchases of goods and services as a measure of intermediate expenditure, tangible fixed assets as a measure of the capital stock, and employment (count of employees) as a measure of the labour input.

First, we deflate revenue, intermediates, and capital using corresponding indices provided by the ONS (for the United Kingdom) and the INSEE (for France) with the base year being 2017.[8] Second, we apply some trimming to the data (Appendix A). Third, we use a second-order polynomial in intermediates, capital and labour to smooth revenue and purge it of measurement error, as suggested in De Loecker et al. [14] and Forlani et al. [15], among others.

Denoting firms by i and time by t the production function we estimate is the following three inputs Cobb–Douglas:

$$R_{it} = L_{it}^{\alpha_L} M_{it}^{\alpha_M} K_{it}^{\alpha_K} A_{it},$$

where A_{it} is Total Factor Productivity (TFP) of the firm i at time t, R_{it} is revenue, L_{it} is labour, M_{it} is intermediate, K_{it} is capital and α_L and α_M and α_K are the related output elasticities. Considering the log production function, we thus have:

$$q_{it} = \alpha_L l_{it} + \alpha_M m_{it} + \alpha_K k_{it} + a_{it}, \tag{1}$$

where lower case letters indicate logs (e.g., $k_{it} = \log K_{it}$). In line with the productivity literature, we assume that the TFP process is driven by an autoregressive component:

$$a_{it} = \phi_a \ a_{it-1} + v_{ait}, \tag{2}$$

where v_{ait} denotes productivity shocks that represent innovations with respect to the information set of the firm in $t-1$ and are iid across firms and time.

In line with the literature, we assume capital k_{it} to be predetermined in the short run, that is the current capital level has been chosen in t-1 and cannot immediately adjust to current period shocks v_{ait}.[9] We further assume, as standard in the literature, that intermediates m_{it} are a variable input free of adjustment costs. This means that intermediates can be optimally chosen in t based on, among others, the particular realisation of v_{ait}. In this respect, we will see later on that intermediates being fully adjustable in the short run allows for a simple rule to pin down the markup of firm i. Concerning labour, we assume it to be a semi-flexible input meaning that it can, to some extent, adjust to current shocks in t but not to the optimal cost-minimising level determined only by wages and marginal productivity.[10]

At time t, firms have already chosen capital and labour and so these inputs are considered as given in their decision process along with the cost of intermediates W_{Mit}. At the same time, productivity a_{it} became known at the time t. We assume firms in t use the above information and constraints to choose intermediates in order to minimise production costs and choose quantity or price (depending upon the features of competition) in order to maximise profits. In this respect, as first highlighted in Hall [17] and further implemented in De Loecker and Warzynski [18], De Loecker et al. [14] and Forlani et al. [15], among others, cost-minimisation of a variable input free of adjustment costs provides a simple rule to pin down markups μ_{it}, which in our Cobb-Douglas production function specification is:

$$\mu_{it} = \frac{\alpha_M}{s_{Mit}}, \tag{3}$$

where s_{Mit} is the share of intermediate expenditure in revenue. Therefore, provided with estimates of the parameters of the production function (1), and in particular of α_M, as well as data on intermediate expenditure and revenue, one can simply compute the firm-specific markup μ_{it} using (3).

In terms of estimating the parameters of the production function (1), we use the approach developed in Wooldridge [19]: (i) we substitute for a_{it} in Equation (1) using (2); (ii) substitute for a_{it-1} using a polynomial in k_{it-1} and m_{it-1}; (iii) in the final augmented production function equation, we do not instrument capital k_{it} but instrument labour and intermediates l_{it} and m_{it} with time lags.[11] We estimate the parameters of the production function separately for each industry while adding as controls a battery of time dummies (and information on foreign ownership for the UK). Standard errors are clustered at the firm level.

In order to provide robustness to our analyses, we also employ a complementary way of estimating TFP. Indeed, Ackerberg et al. [16] and Gandhi et al. [20] raise some concerns over the capacity of the proxy variable approach to identify the parameters of gross-output production functions. In order to allay those concerns, we also estimate TFP using a value-added production function, while still building on the proxy variable approach and the same moment conditions used above. Reassuringly, our key results on the determinants of the UK and French spatial productivity differences across regions are virtually unaffected when using this complementary measure of firm TFP. To provide further robustness, we also perform simple OLS estimations of the production function (1).

5 | Some Data Highlights

While the French data we exploit in this analysis have been used in many other studies in the past,[12] the value of using such data for us lies in its close comparability to the UK data (as shown below), which allows a thorough comparison of the two countries, and in particular allows the granular spatial analysis we perform later on. In contrast, the UK dataset we have constructed has never been assembled before, and this is one of our key contributions. The ARD/ABS surveys administered by the ONS have so far been the workhorse when analysing firm productivity in the United Kingdom. The advantage of our data over the ARD/ABS surveys is the much higher coverage of firms, allowing us, for example to study the spatial distribution of productivity at a more granular level, such as the TTWA level. In what follows, we provide some highlights of both the UK and the French datasets to show evidence of comparability while at the same time pointing to the differences between our UK data and the ARD/ABS surveys.

Tables B1 to B4 in the Supporting Information provide (for both the UK and France) estimates of the parameters of the production function (for each industry) obtained with our instrumental variables approach à la Wooldridge [19] (we label such estimates WLD). Inspection of Tables B1 to B4 reveals that coefficients are quite precisely estimated and have the expected magnitude for a three-input production function, namely an elasticity of intermediates around 0.7–0.8, an elasticity of labour around 0.2, and an elasticity of capital around 0.02–0.05. Furthermore, the under-identification tests and the weak identification F-statistics clearly indicate that our instruments are strong.

For a taste of the size and coverage of our dataset, we provide in Tables 1 and 2 below some key summary statistics across all years. Considering the United Kingdom, our dataset spans over 9,954,131 observations across the time frame 2004–2017. The average firm has £4.3 million revenue,[13] £3.2 million value of intermediates, a £2.4 million capital stock and 22 workers. Standard deviation values are almost two orders of magnitude higher than mean values, indicating that our data cover both very small and very large firms. A closer look at the 5th and 95th percentiles further confirms this. Considering capital, for example the firm in the 5th percentile has a capital stock of just over 1000 pounds, while the 95th percentile firm has a capital stock of about 664,000 pounds, which is still below the mean of 2.4 million pounds, with the latter being driven up by the presence of a few very big firms. As for France, our dataset spans over 17,641,530 observations across the time frame 2000–2017. The average firm has a 3 million euros revenue, a 2.1 million value of intermediates, a 1.8 million capital stock, and 13 workers. Similarly, the data cover both very small and very large firms. Considering capital, for example the firm in the 5th percentile has a capital stock of just over 6000 euros, while the 95th percentile firm has a capital stock of about 1.7 million euros, which is still below the mean of 1.8 million euros, with the latter being driven up by the presence of a few very large firms.

Table B5 in the Supporting Information provides the same information as Table 1 for the United Kingdom while using the ARDx database (a database combining the ARD and ABS surveys) over the same time period. Comparison of the two Tables highlights how the ARDx database has both a much lower coverage than our data (15 times less firms) and a bias towards large firms, which is in line with the sample design of the ARD/ABS surveys covering all the big firms and a small fraction of the medium and small firms [21]. At the same time, the total capital stock from the ARDx database, computed as the sum of the capital stock of each firm-year observation, amounts to about 11.3 trillion pounds for the period 2004–2017 while the corresponding figure for our database is 24.1 trillion pounds, so allaying concerns over the undermeasurement of capital emanating from the FAME database. Last but not least, in our database there are more

TABLE 1 | UK Data: Key summary statistics across all years.

	Mean	St.dev.	p5	p95	N. observ.
Revenue	4305.85	219,289.69	31.93	5287.83	9,954,131
Intermediates	3159.46	171,339.37	7.38	3655.35	9,954,131
Capital	2424.68	245,007.26	1.20	664.60	9,954,131
Employment	21.95	622.40	1	38	9,954,131

Note: Revenue, intermediates, and capital are measured in 1000 pounds. Values have been deflated using indexes provided by the ONS, with the base year being 2017. Employment is the number of employees counted, including the owner(s).

TABLE 2 | French Data: Key summary statistics across all years.

	Mean	St.dev.	p5	p95	N. observ.
Revenue	2968.16	92,725.74	64.62	5751.42	17,641,530
Intermediates	2110.03	70,285.80	23.92	3897.24	17,641,530
Capital	1806.16	174,892.63	6.27	1717.29	17,641,530
Employment	12.74	372.00	1	32	17,641,530
Wage bill	581.54	19,062.60	15.20	1344.75	17,641,530

Note: Revenue, intermediates, capital, and wage bill are measured in 1000 euros. Values have been deflated using indexes provided by the INSEE, with the base year being 2017. Employment is the number of employees.

TABLE 3 | UK Data: Number of firms and total employment covered by year.

Year	Number of firms	Total employment
2004	642,748	13,812,662
2005	681,104	14,198,956
2006	695,050	14,470,623
2007	717,933	14,851,475
2008	701,827	15,378,391
2009	684,485	15,307,760
2010	681,465	15,294,427
2011	700,898	15,544,064
2012	692,865	15,899,287
2013	716,939	16,263,075
2014	728,632	16,362,476
2015	740,365	16,609,343
2016	755,413	17,058,927
2017	814,407	17,441,714

Note: Employment is number of employees count including the owner(s). Data are organised by fiscal year with, for example the year 2017 corresponding to the fiscal year 2017–2018.

than 100 observations for each TTWA in each year, while in the ARDx less than half of TTWAs would meet this criterion. In this respect, we thus believe that our dataset is better suited than the ARDx to, for example analysing firm TFP across a fine spatial scale like the TTWA geographical disaggregation.

Tables 3 and 4 below provide a breakdown of the number of firms (and the related overall employment) in our two datasets by year. For the UK, the number of firms rises from 642,748 in 2004 to 814,407 in 2017. Overall employment covered by our dataset is between 14 and 17 million.[14] Considering France, the number of firms varies by year, with a maximum of 1,144,423 in 2007 and a minimum of 871,200 in 2013. Overall employment covered by our dataset is rather stable across years and above 12 million. Tables B6 and B7 in the Supporting Information contain instead an industry breakdown of the number of firms and related employment for the year 2017.

Finally, Tables 5 and 6 deliver average (employment weighted) apparent labour productivity, labour productivity, OLS TFP, WLD TFP, and markups by year.[15] Considering the United Kingdom, Table 5 indicates that, while total factor productivity (both OLS TFP and WLD TFP) has been both only very lightly and very briefly affected by the financial crisis, the same is not true for markups, apparent labour productivity, and labour productivity, which is consistent with evidence provided in analyses based on the ARD/ABS surveys like Harris and Moffat [11] and Jacob and Mion [12]. Inspection of markups reveals that they recovered their pre-financial crisis level around 2015,[16] while for labour productivity the recovery year is 2016.[17] In this respect, Tables B9 and B10 in the Supporting Information show that results are similar if we split the sample into single-TTWA firms (essentially small and medium firms) and multi-TTWA firms (essentially large firms) with the recovery being stronger for multi-TTWA firms. For single-TTWA firms, labour productivity in 2017 is still below its pre-financial crisis level.

TABLE 4 | French Data: Number of firms and total employment covered by year.

Year	Number of firms	Total employment
2000	1,025,542	12,006,862
2001	1,012,852	12,294,591
2002	1,021,618	12,440,875
2003	1,044,963	12,073,664
2004	1,077,003	12,700,392
2005	1,046,706	12,570,017
2006	1,113,641	12,956,367
2007	1,144,423	13,018,617
2008	927,707	12,636,208
2009	927,597	12,294,506
2010	937,374	12,527,977
2011	936,053	12,659,021
2012	919,392	12,512,977
2013	871,200	12,328,195
2014	909,314	12,383,382
2015	885,391	12,543,022
2016	940,728	12,325,677
2017	900,026	12,406,277

Note: Employment is number of employees.

Moving on to France, Table 6 suggests that it is not entirely clear whether total factor productivity had by 2017 picked up its pre-financial crisis level (OLS vs. WLD). On the other hand, apparent labour productivity and labour productivity have been little affected by the financial crisis.[18] Inspection of markups reveals that they have not yet recovered their pre-financial crisis level, suggesting that firms struggle to achieve pre-financial crisis profit margins. Results are similar if we split the sample into single-TTWA firms and multi-TTWA firms (Tables B11 and B12 in the Supporting Information) with the recovery being stronger for multi-TTWA firms. For multi-TTWA firms, total factor productivity has definitely reattained its pre-financial crisis level.

6 | Regional Productivity Differences: Conceptual Framework

A stylised fact of economic geography is that the productivity of firms increases with city size and urban density [1],[19] and a large literature going back to Marshall [2] explores the question of why cities have this productivity advantage. Micro-foundations put forward for these agglomeration externalities are typically grouped under the headings sharing, matching, learning and sorting [3, 4] and include different forms of knowledge spillovers between firms, costly trade, pro-competitive effects of city size and sorting of workers [5]. The empirical literature suggests a rather consistent, across countries and years, range for the elasticity of productivity with respect to city size. Rosenthal and Strange [6] and Combes and Gobillon [1] provide summaries of this literature and agree on a range for the key elasticity of productivity with respect to density of 0.02–0.10.[20] These findings are robust

TABLE 5 | UK Data. Average (employment weighted) apparent labour productivity, labour productivity, OLS TFP, WLD TFP and markups by year.

Year	Apparent Lab. Prod.	Lab. Prod.	OLS TFP	WLD TFP	Markups	No. of firms
2004	192,796	57,577	3.490	3.036	1.558	642,748
2005	202,646	57,349	3.514	3.055	1.565	681,104
2006	201,485	57,688	3.546	3.084	1.545	695,050
2007	209,504	56,681	3.543	3.079	1.561	717,933
2008	188,056	47,892	3.537	3.070	1.533	701,827
2009	179,307	47,832	3.528	3.062	1.534	684,485
2010	189,490	44,674	3.547	3.075	1.512	681,465
2011	191,634	43,756	3.548	3.074	1.513	700,898
2012	191,446	46,667	3.557	3.084	1.527	692,865
2013	190,029	47,480	3.594	3.123	1.532	716,939
2014	199,459	50,321	3.661	3.193	1.559	728,632
2015	197,796	54,829	3.706	3.237	1.570	740,365
2016	204,431	58,751	3.703	3.233	1.591	755,413
2017	206,930	59,777	3.736	3.268	1.620	814,407

Note: Employment is the number of employees count including the owner(s). Data are organised by fiscal year with, for example the year 2017 corresponding to the fiscal year 2017–18. Revenue, intermediates, and capital have been deflated using indexes provided by the ONS with the base year being 2017. Apparent labour productivity is computed as firm revenue (in 2017 pounds) over firm employment. Labour productivity is computed as firm value added (in 2017 pounds) over firm employment. OLS TFP is firm-level total factor productivity obtained from a 3 inputs (intermediates, labour and capital) Cobb-Douglas production function where revenue is the output measure and coefficients are estimated (separately by SIC industry) using the OLS estimator. WLD TFP is firm-level total factor productivity obtained from a 3 inputs (intermediates, labour and capital) Cobb-Douglas production function where revenue is the output measure and coefficients are estimated (separately by SIC industry) using a method consistent with Wooldridge [19]. Markups are estimates of the firm-level price to marginal cost ratio and are obtained from WLD TFP estimations and the share of intermediates in revenue as developed in De Loecker and Warzynski [18]. All firm-level variables have been aggregated using firm employment as weight.

to the endogeneity of current economic density and in particular to the use of long lags of historical density as instruments for current density [24, 25].

While most geographers would typically consider regions as the unit of analysis and directly work at this level of aggregation, economists are increasingly using firms or even establishments as the unit of analysis around which to reconstruct and attribute differences in economic performance across regions. Crucially, the two approaches do not seem to provide the same magnitudes regarding, for example the elasticity of productivity with respect to local density. More specifically, Jacob and Mion [9] provide evidence for French manufacturing firms highlighting the importance of weighting in going from the firm-level (micro) to the regional-level (macro) productivity. They find smaller values for the elasticity of productivity with respect to population density when using unweighted firm-level regressions, while getting somewhat larger values when considering revenue- or employment-weighted firm-level regressions.

The productivity of a region clearly reflects the productivity of its firms. However, the aggregate productivity of a region is the weighted average (typically by employment) of the productivity of the firms located in the region and not the simple average. When running unweighted regressions, in which firms are the unit of analysis, to measure productivity differences across space one is essentially comparing the average firm across different locations irrespective of the firm size distribution, and its link to productivity, within regions. The link between micro and macro is restored if one runs weighted regressions (as explained better below) and the coefficients from the unweighted and weighted regressions do not need to be the same. One reason they could differ is a varying (across regions) correlation between firm size and firm productivity. For example, if denser regions are characterised by a higher correlation between firm size and firm productivity, unweighted differences in productivity across space will be magnified when weighting. Another reason for differences between coefficients is the heterogeneity (along the productivity dimension) of the elasticity of productivity with respect to density. For example, if more productive firms enjoy disproportionate gains from the density of economic activities, unweighted differences in productivity across space will again be magnified when weighting because (on average) more productive firms are larger.

In what follows we extend the analysis of Jacob and Mion [9] beyond manufacturing to the whole private sector for both France and the United Kingdom while digging into the above-mentioned explanations for the larger values of the elasticity of productivity with respect to population density when using weighted firm-level regressions as compared to unweighted firm-level regressions. We are interested in the variation of TFP across regions and how it is affected by aggregation/weighting. The baseline estimation equation is

$$\overline{a}_{it} = \gamma density_{r(it)} + I_{r(it)} + I_t + \epsilon_{it}, \tag{4}$$

where

- $\overline{a}_{it}$ is (log) WLD TFP demeaned by the corresponding industry average (we net out composition effects)
- $density_{r(it)}$ is the log density of population in region r where firm i is observed at time t,[21]
- $I_{r(it)}$ and I_t are macro region and year dummies
- ϵ_{it} is an error term.

TABLE 6 | French Data. Average (employment weighted) apparent labour productivity, labour productivity, OLS TFP, WLD TFP and markups by year.

Year	Apparent Lab. Prod.	Lab. Prod.	OLS TFP	WLD TFP	Markups	N. of firms
2000	223,361	64,876	1.654	2.487	1.265	1,025,542
2001	224,933	66,008	1.659	2.498	1.256	1,012,852
2002	228,402	66,431	1.651	2.480	1.261	1,021,618
2003	231,482	66,605	1.628	2.332	1.257	1,044,963
2004	229,501	66,766	1.657	2.483	1.266	1,077,003
2005	232,417	67,696	1.663	2.490	1.264	1,046,706
2006	235,324	67,523	1.662	2.486	1.261	1,113,641
2007	239,435	67,855	1.666	2.489	1.262	1,144,423
2008	233,018	66,057	1.557	2.340	1.231	927,707
2009	222,918	67,669	1.614	2.482	1.246	927,597
2010	226,106	66,514	1.606	2.469	1.241	937,374
2011	239,428	66,912	1.617	2.482	1.232	936,053
2012	237,460	67,776	1.620	2.506	1.240	919,392
2013	237,251	68,404	1.626	2.523	1.244	871,200
2014	239,378	68,546	1.621	2.527	1.237	909,314
2015	232,609	67,769	1.622	2.503	1.246	885,391
2016	238,234	69,863	1.621	2.498	1.243	940,728
2017	242,910	69,482	1.620	2.493	1.244	900,026

Note: Employment is the number of employees. Revenue, intermediates and capital have been deflated using indexes provided by the INSEE, with the base year being 2017. Apparent labour productivity is computed as firm revenue (in 2017 euros) over firm employment. Labour productivity is computed as firm value added (in 2017 euros) over firm employment. OLS TFP is firm-level total factor productivity obtained from a 3-inputs (intermediates, labour and capital) Cobb-Douglas production function where revenue is the output measure and coefficients are estimated (separately by NACE rev2 industry) using the OLS estimator. WLD TFP is firm-level total factor productivity obtained from a 3-inputs (intermediates, labour and capital) Cobb-Douglas production function where revenue is the output measure and coefficients are estimated (separately by NACE rev2 industry) using a method consistent with Wooldridge [19]. Markups are estimates of the firm-level price to marginal cost ratio and are obtained from WLD TFP estimations and the share of intermediates in revenue as developed in De Loecker and Warzynski [18]. All firm-level variables have been aggregated using firm employment as weight.

We perform both unweighted and weighted (by the share of employment of the firm i in the total region r employment at the time t) OLS estimations of Equation (4) and cluster standard errors at the region-year (ZE for France and TTWA for the United Kingdom) level. We are interested in the estimates of γ and $I_{r(it)}$ and in particular by how much, if anything, those estimates get larger if we consider weighting when we switch from the micro (firms) to the macro (regions). More formally, consider first aggregating firm productivity $\overline{a}_{it}$ (using our weights) in the region r and year t level and then regressing this average region-year productivity on density, as well as macro region and year dummies, while using robust standard errors. The resulting OLS coefficients and standard errors of this 'aggregate' regression will be, by the properties of OLS, identical to those obtained from weighted OLS estimations of Equation (4) at the firm level with region-year clustering of the standard errors. In light of this, unweighted and weighted regressions of Equation (4) allow us to navigate from the micro firm level to the macro regional level. At the same time, the R2 of the weighted regression of Equation (4) and the aggregate regression will be different because in the latter heterogeneity in productivity across firms within a region-year has been eliminated.

7 | Results

In what follows, we focus on the samples of single TTWA firms (for the UK) and single ZE firms (for France), those firms that we can uniquely associate with a region. Such firms may thus have more than one establishment, but such establishments need to be located in the same TTWA/ZE. In the next section, we provide some robustness results including multi-TTWA/ZE firms in the analysis. The robustness results largely confirm our findings based on single TTWA/ZE firms.

Table 7 provides estimates of Equation (4) for the United Kingdom. The first column contains unweighted estimates, while the second column delivers weighted results. The weighted γ is around 2.1% and so in line with the literature, while the unweighted coefficient of density stands at about 1.8%. At the same time, macro region dummies (the reference category being London) are all negative and strongly significant, while being typically larger in magnitude in the case of weighted regressions. These results suggest some amplification of unweighted differences in productivity across space when considering weighting.[22] For example, our estimates imply that the aggregate productivity difference between the median density region (Banbury, East Midlands) and London is 16.6%, while the unweighted productivity difference between firms in the two regions is 12.4%, so the latter accounts for about 75% of the aggregate difference. Furthermore, both unweighted and weighted estimates point to a substantial productivity gap, over and beyond what can be explained by density, of all UK regions relative to London.

As suggested above, the difference between the two sets of estimates could be driven by (1) a correlation between firm size and productivity varying across regions and, in particular, increasing with density; (2) a productivity return on density being

TABLE 7 | UK: Density regressions.

	Unweighted	Weighted
Log density	0.0178***	0.0208***
	(0.0007)	(0.0009)
	Reference category is London	
East Midlands	−0.0676***	−0.1004***
	(0.0027)	(0.0072)
East of England	−0.0421***	−0.1056***
	(0.0031)	(0.0082)
North East	−0.0765***	−0.1021***
	(0.0030)	(0.0084)
North West	−0.0767***	−0.1252***
	(0.0029)	(0.0074)
Northern Ireland	−0.0211***	−0.0877***
	(0.0032)	(0.0073)
Scotland	−0.0128***	−0.0749***
	(0.0043)	(0.0077)
South East	−0.0293***	−0.0902***
	(0.0033)	(0.0074)
South West	−0.0718***	−0.1129***
	(0.0030)	(0.0069)
Wales	−0.0828***	−0.1373***
	(0.0030)	(0.0077)
West Midlands	−0.0772***	−0.1147***
	(0.0026)	(0.0070)
Yorkshire and The Humber	−0.0804***	−0.1137***
	(0.0028)	(0.0076)
Observations	9,663,658	9,663,658
R^2	0.0094	0.0071
R^2 'aggregate'		0.2613

Note: The dependent variable is firm (log) WLD TFP demeaned by the corresponding industry average. Log density is the 2015 population density of the TTWA r where firm i is located at time t. Year dummies (not reported) and macro region dummies (London being the reference category) are included in the regressions. Column one provides simple OLS regressions of Equation (4) while column two shows weighted OLS regressions of Equation (4) where the weight is the share of the employment of firm i and time t in overall regional employment in year t. R-squared 'aggregate' refers to the R-squared of the equivalent, to the weighted firm-level regression in column two, 'aggregate' regression at the region-year level. Standard errors are clustered by TTWA-year. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

stronger for the most productive firms. In order to analyse the first hypothesis, we compute the correlation between firm size and productivity for each of the 228 TTWAs in the United Kingdom and plot in Figure 1 these correlations against the density of the region. Inspection of Figure 1 reveals that (1) the correlation between firm size and productivity within a region is quite low and sometimes negative across UK TTWAs and (2) the relationship between these correlations and region density is not positive and actually slightly negative (red regression line in the Figure 1). These findings indicate that, if anything, the varying correlation between firm size and productivity across regions should *reduce* and not amplify spatial productivity differences when going from the micro to the macro.

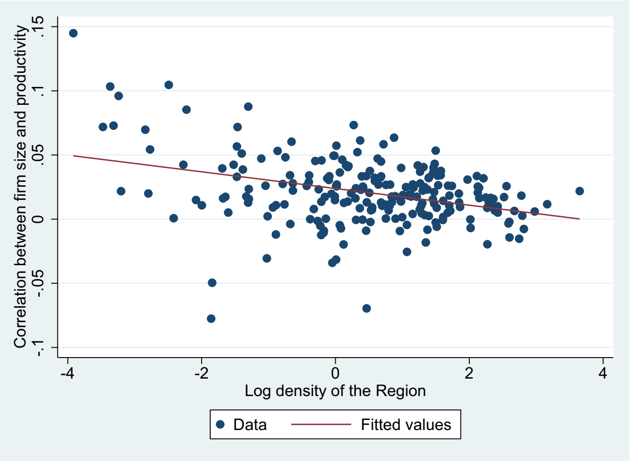

FIGURE 1 | UK: Correlation between firm productivity and size within each region. The figure provides a scatter plot of two variables for each TTWA. The variable on the y-axis in the correlation (across firms and years within a region) between firm employment and (log) WLD TFP demeaned by the corresponding industry average. The variable on the x-axis is the 2015 population density of the TTWA r. The red line indicates the regression line.

With the aim of exploring room for the second explanation, we report in Table 8 the results of quantile estimations (for each decile of the conditional distribution of productivity) of Equation (4) while focusing on the coefficient of density. Table 8 indicates that the productivity return of density is increasing along the deciles of the *conditional* productivity distribution, ranging from 1.1% for the first decile to 3.5% in the 9th decile. This finding is reminiscent of the 'dilating' of the productivity distribution in larger regions found for France in Combes et al. [13], and it is such an increasing productivity return of density that magnifies firm-level productivity differences across the UK space, when going from the micro to the macro, and not a varying correlation between firm size and productivity across space.

In order to demonstrate further that the productivity return of density is increasing along the conditional productivity distribution, we plot in Figure B1 in the Supporting Information the kernel density of the distribution of productivity for both TTWAs with density below the median (low density regions) and TTWAs with density above the median (high density regions). Figure B1 clearly shows how the distribution for high-density regions is not a simple shift of the distribution for low-density regions but rather a dilation, or a stretching, of the latter from the right-hand side of the distribution.

Moving forward, Table 9 provides the equivalent information of Table 7 for France. As can be noticed, the weighted γ is around 2%. This is in line with the literature while the unweighted coefficient of density is much smaller standing at about 0.4%.[23] At the same time, macro region dummies (the reference category being Île-de-France, i.e., Paris) are all negative but small in magnitude and often not significant. On the one hand, these results indicate, contrary to the United Kingdom, the absence of a strong productivity gap (over and beyond what can be attributed to density) between the core region of France and the rest of the country. On the other hand, they also suggest a stronger

TABLE 8 | UK: Quantile regressions and the heterogeneous impact of density across the productivity distribution.

Variables	1st decile	2nd decile	3rd decile	4th decile	5th decile	6th decile	7th decile	8th decile	9th decile
Log density	0.0115***	0.0133***	0.0124***	0.0115***	0.0122***	0.0144***	0.0180***	0.0238***	0.0355***
	(0.0004)	(0.0002)	(0.0002)	(0.0001)	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0004)

Note: The dependent variable is firm (log) WLD TFP demeaned by the corresponding industry average. Log density is the 2015 population density of the TTWA r where firm i is located at time t. Year dummies and macro region dummies (not reported) are included in the regressions. Columns one to nine provides quantile regressions of Equation (4) on deciles one to nine. Robust standard errors provided. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

TABLE 9 | France: Density regressions.

	Unweighted	Weighted
Log density	0.0041***	0.0196***
	(0.0010)	(0.0051)
	Reference category is Île-de-France	
Auvergne-Rhône-Alpes	−0.0098*	−0.0207
	(0.0057)	(0.0185)
Bourgogne-Franche-Comté	−0.0162***	−0.0189
	(0.0059)	(0.0181)
Bretagne	0.0096*	0.0033
	(0.0051)	(0.0221)
Centre-Val de Loire	−0.0118**	−0.0134
	(0.0059)	(0.0185)
Grand Est	−0.0181***	−0.0115
	(0.0051)	(0.0196)
Hauts-de-France	−0.0111**	−0.0178
	(0.0044)	(0.0215)
Normandie	−0.008	−0.0064
	(0.0052)	(0.0215)
Nouvelle-Aquitaine	−0.0229***	−0.0320*
	(0.0055)	(0.0181)
Occitanie	−0.0410***	−0.0367**
	(0.0068)	(0.0184)
Pays de la Loire	−0.0001	−0.001
	(0.0053)	(0.0198)
Provence-Alpes-Côte d'Azur	−0.0375***	−0.031
	(0.0062)	(0.0206)
Multi-region	−0.0182*	−0.0132
	(0.0095)	(0.0226)
Observations	16,595,355	16,595,355
R^2	0.0050	0.0089
R^2 'aggregate'		0.2458

Note: The dependent variable is firm (log) WLD TFP demeaned by the corresponding industry average. Log density is the 2009 population density of the ZE r where firm i is located at time t. Year dummies (not reported) and macro region dummies (Paris being the reference category) are included in the regressions. Column one provides simple OLS regressions of Equation (4) while column two shows weighted OLS regressions of Equation (4) where the weight is the share of the employment of firm i and time t in overall regional employment in year t. R-squared 'aggregate' refers to the R-squared of the equivalent, to the weighted firm-level regression in column two, 'aggregate' regression at the region-year level. Standard errors are clustered by ZE-year. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

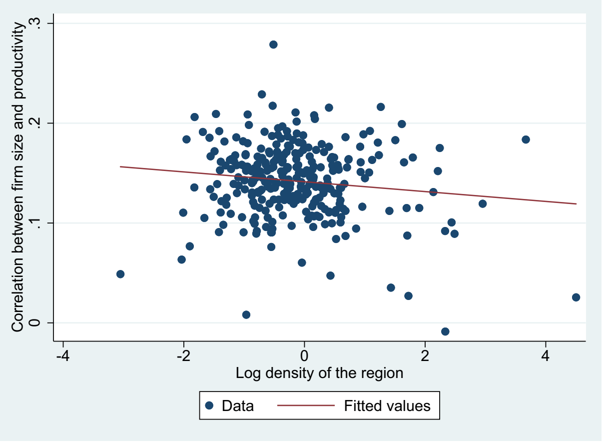

FIGURE 2 | France: Correlation between firm productivity and size within each region. The Figure provides a scatter plot of two variables for each ZE. The variable on the y-axis in the correlation (across firms and years within a region) between firm employment and (log) WLD TFP demeaned by the corresponding industry average. The variable on the x-axis is the 2009 population density of the ZE r. The red line indicates the regression line.

(for France compared to the United Kingdom) amplification of unweighted differences in productivity across space when considering weighting. For example, our estimates imply that the aggregate productivity difference between the median density region (Saint-Dié-des-Vosges, Grand Est) and Paris is 10.5%, while the unweighted productivity difference between firms in the two regions is 3.77%, so the latter accounts for only about 36% of the aggregate difference.

Figure 2 further qualifies our results by showing (as in the case of the UK) that the relationship between the correlation of firm size and productivity within a region and region density is not positive and, if anything, slightly negative (red regression line in the Figure). A comparison of Figures 1 and 2 also reveals that the correlation between firm size and productivity is more on the positive side across French ZEs as compared to UK TTWAs, and this extends to the nationwide correlation between firm employment and productivity standing at 0.0243 for France (0.0712 when using our weights instead of firm employment) and 0.0122 for the UK (−0.0046 when using our weights instead of firm employment). If the UK had the French correlations, aggregate productivity would be higher.[24]

Furthermore, Table 10 shows the results of quantile estimations (for each decile of the conditional distribution of productivity) of Equation (4) for France while focusing on the coefficient of density. As in the case of the United Kingdom, the productivity return of density is increasing along the deciles of the conditional productivity distribution, and this is the key driver of the magnification of productivity differences across space when going from the micro to the macro. Contrary to the United Kingdom, though, the productivity return of density is actually negative for the first few deciles of the productivity distribution signalling an issue that France seems to have about large agglomeration and low productive firms. Interestingly, this finding has no counterpart in the productivity distribution analysis of Combes et al. [13]. At the same time, Figure B2 in the Supporting Information provides, as in the UK case, further evidence of the dilation, or

TABLE 10 | France: Quantile regressions and the heterogeneous impact of density across the productivity distribution.

Variables	1st decile	2nd decile	3rd decile	4th decile	5th decile	6th decile	7th decile	8th decile	9th decile
Log density	−0.0100***	−0.0052***	−0.0018***	0.0007***	0.0030***	0.0050***	0.0075***	0.0113***	0.0189***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)

Note: The dependent variable is firm (log) WLD TFP demeaned by the corresponding industry average. Log density is the 2009 population density of the ZE r where firm i is located at time t. Year dummies and macro region dummies (not reported) are included in the regressions. Columns one to nine provides quantile regressions of Equation (4) on deciles one to nine. Robust standard errors provided. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

right-stretching, of the productivity distribution when comparing ZEs with density below the median (low density regions) and ZEs with density above the median (high density regions).

8 | Robustness

Using both single-region and multi-region firms. As anticipated above, we have produced equivalent results to those presented in Section 6 for the full sample of single-region and multi-region firms. In doing so, we exploit information on the different establishments of a firm and allocate the same productivity to all of the establishments of a firm while using establishment-level employment to weight our regressions. Indeed, the unit of analysis in these robustness exercises, where density regressions are presented in Tables B13 and B14 in the Supporting Information, switches from the firm to the establishment. Moving to our findings, Figures B3 and B4 in the Supporting Information convey a very similar message to Figures 1 and 2, so the relationship between the correlation of establishment size and productivity within a region and region density is not positive and, if anything, slightly negative. At the same time, French ZEs are characterised by a more positive correlation between establishment size and productivity than UK TTWAs. Furthermore, Tables B15 and B16 in the Supporting Information indicate that the productivity return of density is (overall) increasing along the establishment productivity deciles. Still, France features some negative productivity returns on density for the first deciles, as in the case of Table 10.

Using a value added production function. We have already discussed above how some scholars in the productivity literature question the capacity of the proxy variable approach to deal with the estimation of production functions in revenue form, like the one we use in our analysis [16, 20]. In order to allay such concerns, we have first estimated productivity using a production function in value added form (see details above) and subsequently used such a complementary productivity measure (that we label WLD-VA TFP) to perform our spatial analysis. Density regressions using such a measure are provided in Tables B17 and B18 in the Supporting Information for the UK and France, respectively. Such Tables confirm the magnification effect related to weighting while overall suggesting a higher productivity return of density. Tables B19 and B20, as well as Figures B5 and B6, overall portray a picture very similar to our baseline with the productivity return of density increasing along the productivity distribution.

Instrumenting current density with historical values. In the typical urban model, density is an endogenous variable whose equilibrium level depends on the fundamentals of both the considered location and the overall interaction between locations. In order to deal with this issue, urban scholars have suggested using deeply lagged density values as instruments for current density, with the idea being that historical density was driven by factors other than those driving current density, and in particular by forces shaping geography prior to the industrial and service revolutions [24, 25]. While the literature finds that the endogeneity of density is often a second-order problem [1], we nevertheless provide in the Supporting Information full analysis based on instrumenting current density with historical values.[25] Results are provided in Tables B21 to B24, as well as in Figures B7 and B8 in the Supporting Information, and strongly confirm our baseline findings.

Using revenue weights instead of employment weights. Throughout our analysis, we use information on firm (or establishment) employment to weight observations because it is the most common practice among statistical institutes like the UK ONS and the French INSEE. However, Melitz and Polanec [28] suggest using revenue to weigh gross-output based TFP. We accomplish this in Tables B25 and B26, where we provide density regressions based on firm-level revenue weighting, confirming previous results. In this respect, we note that the finding that the productivity return of density is increasing along the productivity distribution is not related to the type of weighting, and so results in Tables 8 and 10 still apply. As for the correlation between firm size (now measured in terms of revenue) and productivity within each location, results are reported in Figures B9 and B10 in the Supporting Information. As the reader can appreciate, the negative relationship between the above correlation and local density still applies to both the UK and France.

Looking at industry patterns: manufacturing and services. Our goal is to document economy-wide patterns for the relationship between firm productivity and density, while highlighting the issue of aggregation and comparing two similar countries like the United Kingdom and France. This is, for example the reason why we (among others) consider firm productivity deviations from the industry average. Having said that, we believe there is a certain interest in looking at industry patterns. While a full analysis is beyond the scope of this paper, we nevertheless report here some results at the industry level. In this respect, the literature highlights the potential differences between manufacturing and services. For example, Berlingieri et al. [29] highlight that for France the correlation between employment and labour productivity is positive for manufacturing, but there is no relation for services, while the returns to density have been found to be stronger in services [30]. Tables B27 to B30 provide density regressions for the United Kingdom and France, while focusing on manufacturing and services.[26] The overall picture emerging from such Tables is one in which the patterns seen above apply also within manufacturing, as well as within services.

9 | Conclusions

We propose a new data resource that attempts to overcome limitations of standard firm-level datasets for the United Kingdom (like the ARD/ABS) by building on administrative data covering the population of UK firms. More specifically, we merge the BSD, VAT and FAME datasets and create a common firm definition encompassing the different firm identifiers used in the three datasets. This delivers us with enough information to estimate TFP (and markups) for an unprecedentedly large number of firms allowing for comprehensive longitudinal analyses and granular regional-level investigations. We also construct a similar dataset for France and use both datasets to (1) provide some highlights of the data and an overall picture of the evolution of aggregate UK and French productivity and markups: (2) analyse the spatial distribution of productivity in both countries at a very fine level of detail – 228 Travel to Work Areas (TTWAs) for the UK and 297 Zones d'emploi (ZEs) for France – while focusing on the role of economic density.

Considering the United Kingdom, while total factor productivity has been both only very lightly and very briefly affected by the financial crisis, the same is not true for markups, apparent labour productivity (revenue per worker) and labour productivity (value added per worker). Inspection of markups reveals that they recovered their pre-financial crisis level around 2015 while for labour productivity the recovery year is 2016. As for France, it is not entirely clear whether total factor productivity has by 2017 reattained its pre-financial crisis level. On the other hand, apparent labour productivity and labour productivity have been little affected by the financial crisis. Inspection of markups reveals that they have not yet recovered their pre-financial crisis level suggesting that French firms struggle to achieve pre-financial crisis profit margins.

In terms of spatial analysis we obtain the following results. First, for both France and the United Kingdom, we find a larger productivity return to density when weighting observations by employment as compared to unweighted regressions. Digging deeper into this reveals, in both cases, that (1) the correlation between firm size and productivity within a region is quite low (and sometimes negative) across regions particularly for the United Kingdom; (2) the relationship between these correlations and region density is not positive and actually slightly negative. These findings indicate that, if anything, the varying correlation between firm size and productivity across regions should *reduce* and not amplify spatial productivity differences when going from the micro to the macro. On the other hand, in both cases, we find evidence that the productivity return of density is increasing along the deciles of the productivity distribution. This finding is reminiscent of the 'dilating' of the productivity distribution in larger regions found for France in Combes et al. [13] and it is such an increasing productivity return of density that magnifies firm-level productivity differences for the United Kingdom and France, when going from the micro to the macro, and not a varying correlation between firm size and productivity across space.

Concerning the comparison between the United Kingdom and France, we document a number of striking differences. In the United Kingdom, the correlation between firm size and productivity within a region is low (compared to France) and sometimes negative. Also the United Kingdom has a problem of productivity being quite unequal across space beyond what could be explained by variations in density (big London gap while little Paris gap). The problem with France is instead the negative productivity return of density for the least productive firms, that is denser places in France nurture too many low productive firms and this creates a big divide between the un-weighted and weighted productivity return of density.

Moving to directions for future research we look forward to seeing more studies, especially studies covering countries other than France and the United Kingdom, tackling the issue of aggregation in measuring the return to density. Indeed, despite a number of common features between France and the United Kingdom, our analysis also reveals important differences so highlighting the importance of country specificities.

Acknowledgements

We are grateful to Josh Martin, Mary O'Mahony, Nick Oulton, Rebecca Riley, Bart Van ark, Tony Venables as well as participants of various seminars and conferences for helpful comments. The authors acknowledge financial support from Labex MME-DII (France) and the Productivity Institute (UK). HM Revenue & Customs (HMRC) agrees that the figures and descriptions of results in the attached document may be published. This does not imply HMRC's acceptance of the validity of the methods used to obtain these figures, or of any analysis of the results. Please note that all statistical results remain Crown Copyright, and should be acknowledged either as such and/or as 'Source: Calculations based on HMRC administrative datasets'. Copyright of the statistical results may not be assigned. Written work intended for publication should include a note to the effect that: 'This work contains statistical data from HMRC which is Crown Copyright. The research datasets used may not exactly reproduce HMRC aggregates. The use of HMRC statistical data in this work does not imply the endorsement of HMRC in relation to the interpretation or analysis of the information'.

Endnotes

[1] See also Combes et al. [3], Mion and Naticchioni [7] and De La Roca and Puga [8] for estimates of the elasticity of worker-level wages with respect to density.

[2] See Bauer and Boussard [10] for a recent study using similar data.

[3] See Appendix A for further details.

[4] This latter restriction allows us to deal with an otherwise inexplicable drop in the number of firms around 2010.

[5] We also checked for consistency and meaning of values across years and kept only firms for which values of sales and acquisitions are both non-missing and greater than zero.

[6] As with the other datasets, we applied some cleaning to the industry classification (which is consistently SIC 2003 in the dataset) and eliminated firms involved in financial and insurance activities.

[7] In our analysis we do not consider French overseas territories (DOMs) as well as Corsica. In order to go from municipalities to ZEs for the whole of our sample period we use a correspondence table provided by the INSEE. The match between municipalities and ZEs works quite well and requires only minor adjustments.

[8] See Appendix A for further details.

[9] Intuitively, the restriction behind this assumption is that it takes a full period for new capital to be ordered, delivered, and installed. Note this

means that k_{it} is uncorrelated with current period shocks v_{ait}. However, this does not mean that k_{it} is uncorrelated with the current productivity *level* a_{it}. For example, investment decisions in $t-1$ are likely to be determined by both the level of capital in $t-1$ and the level of productivity in $t-1$. In this light, k_{it} should be correlated with a_{it-1} and so with a_{it}. See Ackerberg et al. [16] for more details.

[10] In sum, l_{it} should be correlated (like intermediates m_{it}) with shocks v_{ait} but the amount of labour in t does not simply reflect wages and marginal productivity implying that it cannot be used to recover markups. As far as the timing is concerned, we assume l_{it} is chosen by firm i at time $t-b$ $(0 < b < 1)$, after k_{it} being chosen in $t-1$ but prior to m_{it} being chosen in t.

[11] We use l_{it-1}, l_{it-2}, m_{it-2} and k_{it-2}.

[12] See Bauer and Boussard [10] for a recent study using similar data.

[13] We have checked the correlation between revenue from VAT data and revenue coming from the BSD and Fame. The correlation between revenue from VAT and revenue from the BSD is quite poor standing at about 0.4. Visual inspection reveals that turnover from the BSD is frequently made up of round numbers that are often not updated across time within a firm. The correlation between turnover from VAT and turnover for FAME (the latter being available for the medium and large firms only) is better and around 0.5 to 0.6. Again, there seems to be a good amount of rounding up in FAME turnover figures.

[14] Firms involved in financial and insurance activities are excluded from our dataset and account for around 1.2 million workers as reported by the ONS official figures. In the ARDx database the overall employment covered is between 10 and 12 million.

[15] In most of our analysis we use employment to weight firm productivity measures while providing below some robustness using revenue weights. Our analysis does not incorporate the insights developed in Baqaee and Farhi [22] to which we refer the reader for further details.

[16] Our findings for markups are not incompatible with those obtained by Black [23]. There are, first of all, some obvious differences in the two analyses. For example, Black [23] uses the ABS/ABI, while we use a more comprehensive database, and so our results also cover those medium and small firms which are missing from the ABS/ABI. On the other hand, Black [23] embraces a longer time span (1997–2019) in which the financial crisis episode is potentially dwarfed by time trends. Finally, some of the markup measures developed in Black [23] also point to a fall in markups, followed by a recovery, around the financial crisis.

[17] Table B8 in the Supporting Information provides complementary information for the UK about the evolution of apparent labour productivity and labour productivity. More specifically, in Table B8 we still provide employment-weighted apparent labour productivity and labour productivity but rather than deflating current values we simply use those current values to compute averages.

[18] Interestingly, when considering apparent labour productivity and labour productivity (the most comparable productivity measures between the two countries) while taking the last year of the data, the UK does not appear to be much less productive than France, which is in contrast to some macro comparisons suggesting that the UK is behind France in terms of productivity. This is of course confined to our data and to those more structured firms employing at least one worker.

[19] See Combes and Gobillon [1] for more explanations about what is, or what is not, captured in these agglomeration economies using TFP as an outcome (e.g., local input costs).

[20] See also Combes et al. [3], Mion and Naticchioni [7] and De La Roca and Puga [8] for estimates of the elasticity of worker-level wages with respect to density.

[21] We use population density by TTWA in 2015 for the UK and population density by ZE in 2009 for France.

[22] At the bottom of the second column of Table 7 we report the R-squared of the equivalent, to our weighted firm-level regression, 'aggregate' regression. Such R-squared is much higher than the one based on the weighted firm-level regression (0.2613 vs. 0.0071) and in line with previous findings in the literature [1].

[23] Although the unweighted coefficient might seem small compared to the literature, the 2.5% coefficient reported in Combes et al. [13] for France actually refers to a weighted regression.

[24] It is not possible to directly measure how much higher UK productivity would be from our analysis unless one is willing to make some strong assumptions. The weighted mean of firm (log) productivity is nothing else than the mean of the product of two variables: the firm-level productivities and the weights. For a given mean and variance of productivities and weights, a 0.0X increase in the correlation between the two variables implies an X% increase in the weighted productivity. Having said that, it is likely that a policy aiming at increasing such correlation would also impact the distributions of firm-level productivity and the weights so complicating the situation.

[25] For the UK we instrument 2015 population density at the TTWA-level with historical TTWA density (reconstructed from census registration district-level data) for the years 1851, 1861 and 1871. The data cover England and Wales [26], as well as Scotland [27], but are not available for Northern Ireland. For France we instrument 2009 population density at the ZE-level with historical ZE density (reconstructed from municipality-level data) in 1831, 1861 and 1891 [9].

[26] Manufacturing comprises NACE rev2 industry codes 10 to 33 while services includes industry codes 45 to 82.

References

1. P.-P. Combes and L. Gobillon, "The Empirics of Agglomeration Economies," in *Handbook of Regional and Urban Economics, 5 of Handbook of Regional and Urban Economics*, ed. G. Duranton, J. V. Henderson, and W. C. Strange (Elsevier, 2015), 247–348.

2. A. Marshall, *The Principles of Economics* (McMaster University Archive for the History of Economic Thought, 1890).

3. P.-P. Combes, G. Duranton, and L. Gobillon, "Spatial Wage Disparities: Sorting Matters!," *Journal of Urban Economics* 63, no. 2 (2008): 723–742.

4. G. Duranton and D. Puga, "Chapter 48 - Micro-Foundations of Urban Agglomeration Economies," in *Cities and Geography, Volume 4 of Handbook of Regional and Urban Economics*, ed. J. V. Henderson and J.-F. Thisse (Elsevier, 2004), 2063–2117.

5. C. Syverson, "What Determines Productivity?," *Journal of Economic Literature* 49, no. 2 (2011): 326–365.

6. S. S. Rosenthal and W. C. Strange, "Evidence on the Nature and Sources of Agglomeration Economies," in *Cities and Geography*. Handbook of Regional and Urban Economics, vol. 4, ed. J. V. Henderson and J.-F. Thisse (Elsevier, 2004), 2119–2171.

7. G. Mion and P. Naticchioni, "The Spatial Sorting and Matching of Skills and Firms," *Canadian Journal of Economics/Revue Canadienne D'économique* 42, no. 1 (2009): 28–55.

8. J. De La Roca and D. Puga, "Learning by Working in Big Cities," *Review of Economic Studies* 84, no. 1 (2017): 106–142.

9. N. Jacob and G. Mion, "On the Productivity Advantage of Cities," *Journal of Economic Geography* 24, no. 5 (2024): 679–705.

10. A. Bauer and J. Boussard, *Market Power and Labor Share* (Economie et Statistique / Economics and Statistics, 2020), 520–521.

11. R. Harris and J. Moffat, "The UK Productivity Puzzle, 2008 to 2012: Evidence Using Plant-Level Estimates of Total Factor Productivity," *Oxford Economic Papers* 69, no. 3 (2017): 529–549.

12. N. Jacob and G. Mion, "The UK'S Great Demand and Supply Recession," *Oxford Bulletin of Economics and Statistics* 85, no. 5 (2023): 993–1022.

13. P.-P. Combes, G. Duranton, L. Gobillon, D. Puga, and S. Roux, "The Productivity Advantages of Large Cities: Distinguishing Agglomeration From Firm Selection," *Econometrica* 80, no. 6 (2012): 2543–2594.

14. J. De Loecker, P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik, "Prices, Markups, and Trade Reform," *Econometrica* 84, no. 2 (2016): 445–510.

15. E. Forlani, R. Martin, G. Mion, and M. Muûls, "Unraveling Firms: Demand, Productivity and Markups Heterogeneity," *Economic Journal* 133, no. 654 (2023): 2251–2302.

16. D. A. Ackerberg, K. Caves, and G. Frazer, "Identification Properties of Recent Production Function Estimators," *Econometrica* 83, no. 6 (2015): 2411–2451.

17. R. E. Hall, "Market Structure and Macroeconomic Fluctuations," *Brookings Papers on Economic Activity* 1986, no. 2 (1986): 285–338.

18. J. De Loecker and F. Warzynski, "Markups and Firm-Level Export Status," *American Economic Review* 102, no. 6 (2012): 2437–2471.

19. J. M. Wooldridge, "On Estimating Firm-Level Production Functions Using Proxy Variables to Control for Unobservables," *Economics Letters* 104, no. 3 (2009): 112–114.

20. A. Gandhi, S. Navarro, and D. A. Rivers, "On the Identification of Gross Output Production Functions," *Journal of Political Economy* 128, no. 8 (2020): 2973–3016.

21. Office for National Statistics, *Annual Respondents Database X, 1997–2020: Secure Access*, 5th ed. (Office for National Statistics, 2023), https://doi.org/10.5255/UKDA-SN-7989-5.

22. D. R. Baqaee and E. Farhi, "Productivity and Misallocation in General Equilibrium," *Quarterly Journal of Economics* 135, no. 1 (2020): 105–163.

23. R. Black, "Estimates of Markups, Market Power and Business Dynamism From the Annual Business Survey, Great Britain: 1997 to 2019," ONS Paper (2022).

24. A. Ciccone, "Agglomeration Effects in Europe," *European Economic Review* 46, no. 2 (2002): 213–227.

25. A. Ciccone and R. E. Hall, "Productivity and the Density of Economic Activity," *American Economic Review* 86, no. 1 (1996): 54–70.

26. H. Jaadla, E. Garrett, A. Reid, and K. Schurer, "Consistent Registration Districts, England and Wales (1851–1911)," 2023, https://doi.org/10.17863/CAM.94859.

27. E. Garrett, "Scotland History of Registration Districts, 1851-1901," 2023, https://doi.org/10.17863/CAM.94394.

28. M. J. Melitz and S. Polanec, "Dynamic Olley-Pakes Productivity Decomposition With Entry and Exit," *RAND Journal of Economics* 46, no. 2 (2015): 362–375.

29. G. Berlingieri, S. Calligaris, and C. Criscuolo, "The Productivity-Wage Premium: Does Size Still Matter in a Service Economy?," *AEA Papers and Proceedings* 108 (2018): 328–333.

30. A. Grover, S. Lall, and J. Timmis, "Agglomeration Economies in Developing Countries: A Meta-Analysis," *Regional Science and Urban Economics* 101 (2023): 103901.

Supporting Information

Additional supporting information can be found online in the Supporting Information section. **Data S1** Supporting Information.

Appendix A

Additional Details on the Data

Details about the construction of the UK database. In the BSD database, the SIC 2007 industry affiliation is not available in the 2004 and 2005 vintages of the data (only SIC 2003 is available), but we exploit the fact that both SIC 2003 and SIC 2007 are available from 2006 onwards to build a correspondence table that we applied to earlier years.

TABLE A1 | Example of correspondences.

BSD firm id	VAT firm id
A	1
A	2
A	3
A	4
A	5
B	3
B	4
B	6
B	7
C	8
D	8

When we match the BSD with the VAT data, we keep only firms present in both datasets. This entails a drop of about 4 to 5 million employees per year, that is concentrated in sectors where public employment is more prevalent. We then match FAME, which at this stage entails a minimal loss in terms of firms in the match, and apply some final cleaning and polishing to the capital stock variable to increase coverage. Finally, we define industries as two digit SIC 2007 codes and apply some grouping in preparation for TFP estimations.

In order to match the different datasets we had to solve the issue of the many-to-many correspondence of firm identifiers. A simple example highlighting the many-to-many issue, and how we deal with it, is reported in Table A1 below.

The example in Table A1 is related to the correspondence between the BSD firm id (for which we use letters) and the VAT firm id (for which we use numbers). Table A1 indicates that the BSD firm id A is linked to many VAT firm id and in particular to 1, 2, 3, 4 and 5. This would not be a problem (being simply a case of one to many) if VAT firm ids 3 and 4 were not also linked to the BSD firm id B, which is also connected to VAT firm id 6 and 7. On the other hand, the case of BSD firm id C and D is simpler because they are both related to the VAT firm id 8, which in turn is not related to other BSD firm id (a simple case of many to one). For our analyses we have devised a looping code that would 'aggregate' BSD and VAT codes in such a way to get, in the case of Table A1, two 'combined firm id' (for which we use Greek letters). The first combined firm id α would correspond to BSD firm id A and B as well as to VAT firm id 1, 2, 3, 4, 5, 6, 7. The second combined firm id β would correspond to BSD firm id C and D as well as to the VAT firm id 8. Once resolved the issue of the many to many cases for the BSD firm id and the VAT firm id, we apply the same procedure using the correspondence between the combined firm id and the FAME firm id, which will generate yet another, more aggregate firm id, encompassing the three different firm identifiers, which we refer to as the 'final firm id'. At the end of the procedure, each original firm id (BSD, VAT and FAME) will be associated to a unique final firm id.

Armed with this notion, we then aggregate the information coming from the three datasets at the final firm id level. For example, we sum the sales of the different VAT codes corresponding to a given final firm id and impute as SIC 2007 code of a final firm id the SIC 2007 code corresponding to the BSD firm id with the largest employment among the different BSD firm ids linked to the final firm id considered.

In order to go from postcodes to 2011 TTWAs for the whole of our sample period we use a postcode directory provided by the ONS. The match between the postcode directory and the postcodes in the data works very well and requires only minor adjustments. Starting from the year 2017 (fiscal year 2017–2018), only the first part of the postcode is available in the BSD data, but fortunately, information on the corresponding TTWA 2011 version is also provided.

Details about the construction of the French database. For both FICUS and FARE we apply the following cleaning to the data. First, we discard observations without information on the municipality where the firm is located and/or with missing SIREN code or industry affiliation. Second, we perform some cleaning on the business start year and the number of employees variables. Third, we use a correspondence table between the NACE rev1 and NACE rev2 to consistently obtain information on the NACE rev2 affiliation of firms for the whole period 2000–2017. Finally, we eliminate firms involved in financial and insurance activities and restrict the sample to firms with at least one employee.

Details about deflation. For France, we use the same industry deflators (base year 2017) for revenue and intermediates while using specific deflators for capital. For the United Kingdom, we employ a double deflation method to deflate all monetary variables to constant 2017 prices. We construct output, intermediates and capital deflators from series provided by the ONS. For the output price deflators, we use the ONS, 'Experimental Industry Level Deflators'. These are available at the 2-, 3- and 4-digit Industry level. They are produced by aggregating industry product deflators based on their use of products in line with the National Accounts supply-use framework. Where the deflators are not available at the 2-digit level, we average the 3- or 4-digit deflators to construct them. To construct deflators for intermediate inputs we make use of the Supply and Use Tables (SUTs) produced by the ONS. The SUTs for 1997 to 2020 are consistent with the UK National Accounts in Blue Book 2022. We use the industries' intermediate consumption values in 2010 to create weights by dividing each 2-digit industry's demand by the total demand for each 2-digit industry. This generates a Leontief matrix which we use as weights to derive input deflators for each 2-digit industry from the output deflators. Capital stock deflators are constructed from the ONS series of annual estimates of net capital stocks and consumption of fixed capital in the United Kingdom, which is provided by asset and sector. These are available in both current prices and chained volume measures. We construct the deflators at the 2-digit industry level. Since tangible fixed assets from FAME are provided as a net value, we obtain the corresponding deflator by dividing the current prices by the chained volume measures of net capital stock at the 2-digit industry level.

Details about trimming in TFP estimations. More specifically, we discard observations where the value of intermediates is higher than the value of sales and further apply a bottom and top trimming of 0.5% (by industry) based on the ratios of: i) intermediates to sales; ii) capital to labour; iii) revenue to labour. Post-TFP estimations, we also discard those (very few) observations with markups below 0.6 and above 20.