Firm-level Productivity Differences: Insights from the OECD's MultiProd Project

Giuseppe Berlingieri
OECD and ESSEC

Patrick Blanchenay
University of Toronto

Sara Calligaris
OECD

Chiara Criscuolo
OECD[1]

ABSTRACT

Productivity plays a central role in shaping the welfare of societies and the competitiveness of countries. Productivity differences, for instance, explain a large share of the differences in income per capita across countries. This article investigates the role of productivity heterogeneity across 18 countries over the period 2001-2012. In particular, it analyses the evidence that emerges from the distributed micro-data approach carried out in the OECD MultiProd project. The main outcome of the project is a unique dataset of harmonized cross-country moments that are representative for the population of firms and comparable across countries even at a detailed industry level. We look at the 90-10 percentile ratio of labour productivity and multifactor productivity and show that: i) productivity dispersion is high in both manufacturing and non-financial market services; ii) it has increased over time, especially in services; iii) a substantial part of this dispersion comes from differences among firms within the same sector of activity in each country; iv) this within-sector dispersion remains the most important component of the overall dispersion for the entire period.

One of the main objectives of economic research is to understand why some nations are more developed than others. A simple measure of economic development, output per capita, illustrates the large disparity found across countries. These disparities largely reflect different levels of productivity across countries. Hall and Jones (1999), for example, find that in the United States output per worker, a measure of labour productivity (LP, henceforth), is 35 times greater than in Niger. LP differences can be partially explained by differences in physical and human capital (Caselli, 1999). However, the main reason for the disparity between these two

1 This article has benefited from helpful feedback from Dirk Pilat, Nick Johnstone, Giuseppe Nicoletti, and CIIE delegates at the OECD. We would also like to thank the editor Andrew Sharpe and two anonymous referees for useful comments and suggestions. All errors remain our own.

Emails: giuseppe.berlingieri@oecd.org, patrick.blanchenay@utoronto.ca, sara.calligaris@oecd.org, and chiara.criscuolo@oecd.org

countries comes from differences in aggregate multifactor productivity (MFP, henceforth), which reflects the overall efficiency with which inputs are combined in the production process. More generally, Prescott (1998) suggests that differences in aggregate MFP are the main driver of international income differences found both across countries and over time. This point is further illustrated by Klenow and Rodiguez-Clare (1997), who conduct a cross-country analysis of 98 nations and suggest that 90 per cent of the country divergence in growth of output per worker correspond to disparity in MFP growth.[2]

In turn, aggregate productivity growth depends closely on MFP at the firm level. If firms increase the efficiency with which they turn inputs into outputs, they can contribute to overall efficiency gains. However, empirical evidence finds substantial heterogeneity in MFP across firms, even within narrowly defined industries. In the US manufacturing sector, the MFP ratio between an industry's 90th and 10th percentile plants is on average 1.92, implying that plants in the 90th percentile roughly make twice the amount with the same inputs as those plants in the 10th percentile (Syverson, 2004). Such dispersion is not only found in developed countries but also in developing ones. For instance, Hsieh and Klenow (2009) find that the ratio of MFP between 90th and 10th percentiles in the manufacturing industries of China and India is on average more than 5:1.

In light of the large dispersion of firms' MFP, analysing industry average productivity does not offer the complete picture: countries might display the same average but very different underlying distributions. This fact has important policy implications. For instance, low average productivity can be explained by either too few firms at the top (lack of innovation), or too many firms at the bottom (weak market selection), two different situations that would entail very different policies.[3] To better design policy strategies, it is therefore essential to understand how firm-level productivity patterns translate into aggregate productivity.

Economists typically attribute differences in MFP across firms to either slow technology adoption or inefficient technology usage.[4] In addition, a growing body of literature attributes high aggregate MFP not only to the efficiency of technology use and speed of adoption but also to the efficient allocation of resources across firms. Resource reallocation can raise aggregate productivity when there is a flow of inputs from low- to high-productivity firms. Conversely, when factors are allocated in an inefficient manner, aggregate productivity is adversely affected.

These important issues have been investigated by two intrinsically interrelated branches of the literature. The reallocation literature typically focuses on the drivers of resource reallocation, such as creative destruction, and upscaling and downscaling of firms, together with the factors that may influence them, such as technological change, regulation and recessions. The misallocation literature typically identifies a specific distortion or a bundle of distortions (policies and/or institutions) and examines the extent to which they adversely impact aggregate productivity. The results obtained in the misallocation literature show that distortions in the economy can have a quantitatively important effect on aggregate productivity.[5]

The OECD contributes to this debate by providing new policy relevant evidence through the

2 See Hsieh and Klenow (2010) and Hopenhayn (2014) for a more recent review of the literature.

3 On this topic see, for example, Malerba and Orsenigo (1995 and 1996), Breschi *et al*. (2000) and Van Dijk (2000).

4 See Parente and Prescott (1994), Comin and Hobijin (2006), Schmitz (2005) and Bloom *et al*. (2013).

5 See Restuccia and Rogerson (2013) for a review on the misallocation literature, and Foster *et al*. (2001, 2002, 2014) for examples of works in the reallocation literature.

multifactor productivity (MultiProd) project. This project provides a comprehensive picture of productivity patterns across a large set of countries over the last two decades.[6] In particular, the project collects micro-aggregated data and moments of the productivity distribution that allow for a cross-country analysis of a wide variety of topics, including productivity heterogeneity, allocative efficiency, misallocation, aggregate productivity growth, and the link between productivity and wages. A similar approach has been used in the past in academic circles (see, for example, Bartelsman, Scarpetta, and Schivardi, 2005; Bartelsman, Haltiwanger, and Scarpetta, 2009), as well as within the OECD (OECD, 2003), the World Bank and, more recently, the European Central Bank.

One of the main contributions of the MultiProd project is to build cross-country harmonized micro-aggregated data of paramount importance for understanding differences in productivity performance across countries. The project relies on a distributed microdata methodology, and the micro-aggregated results (at the cell level) are collected, checked and analysed at the OECD.[7]

An important aspect of the methodology is to make sure the data are comparable across countries. Therefore, productivity is measured in exactly the same way across countries, conditional on the available data. To further ensure harmonization and representativeness, in particular for countries where MultiProd relies on production surveys, an appropriate set of weights is built using information from business registers, which typically cover the whole population of firms. We use these weights to reweight production surveys.

Many studies based on micro-level datasets adopt a resampling procedure in order to achieve a representative dataset (to name a few, for example, Schwellnus and Arnold, 2008, and Arnold *et al.*, 2008). However, MultiProd is to our knowledge the first project based on a distributed microdata approach to have implemented a highly disaggregated, variable-specific, reweighting strategy for both representativeness and aggregation. This reweighting strategy allows us to compute moments representative for the population of businesses and suitable for cross-country comparison even at the two-digit industry or at a more disaggregated level.

The output of the algorithm is a collection of statistics at different levels of aggregation over the 1994-2012 period, depending on data availability. It allows for various decompositions of aggregate productivity level, growth and dispersion to understand the role of particular industries or groups of firms in explaining aggregate outcomes (e.g. small vs large; multinational corporations; old vs young; low vs high productivity; etc.). For instance, changes in the overall productivity dispersion are decomposed to quantify how much of an increase in dispersion is due to an increase in within-industry dispersion and how much comes from a reallocation of resources to industries characterized by a higher dispersion. Moreover, the role of the largest firms can be investigated in great detail and compared to that of the most productive firms ("frontier firms"). Finally, MultiProd attempts to shed light on the nature of wage inequality

6 The time period is to some extent country specific depending on data availability and is limited to more recent years for some countries. For more details on MultiProd, see Berlingieri *et al.,* (2017a).

7 The OECD pioneered this methodology at the beginning of the 2000s (OECD, 2003). It currently follows the distributed microdata approach in three ongoing projects: MultiProd (Multifactor Productivity), DynEmp (Dynamics of Employment) and MicroBeRD (Microdata-based Analysis of Business Expenditure on R&D). See further details in Section 2.

across countries, as well as on the effects of various policies (e.g. employee protection legislation, minimum wage, coordination in wage setting) on the dispersion of wages and productivity.

This article focuses on one specific pillar of the MultiProd project: productivity heterogeneity. It describes the main methodological tools used to carry out the analysis on productivity heterogeneity, and the specific contributions of the MultiProd project. After a description of the distributed microdata approach and of the MultiProd dataset, we present some evidence on productivity heterogeneity looking at the 90-10 percentile ratio of LP and MFP. We show a twofold result: i) dispersion is high in both manufacturing and non-financial market services; ii) it has increased over time, especially in services. Furthermore, we decompose the aggregate dispersion of productivity into within-industry and between-industry components: the within-industry dispersion accounts for most of the total dispersion in both manufacturing and non-financial market services. With a share of more than 80 per cent in almost all years, the within sector variance of productivity is the most important component for the entire period. However, the pattern over time displays a constant or increasing trend until 2008, when it is reversed into a slight decline during the Great Recession.

The rest of the article is organized as follow: the next section provides an overview of the data and methodology used, and the subsequent discusses the MultiProd dataset. Section "Productivity Heterogeneity" looks at the evolution of productivity dispersion across sectors and over time, as well as at the decomposition of the productivity dispersion across sectors. The last section concludes.

Data and Methodology

The distributed microdata approach

In recent years, the policy and research communities' interest in harmonized cross-country microdata has increased significantly. This has been partly driven by improvements in computing power but, fundamentally, it reflects the recognition that microdata are instrumental for understanding the growing complexity in the way economies work and the heterogeneity in economic outcomes.

While considerable progress has been made in providing researchers with secure access to official micro-data on firms at the level of a country, significant obstacles remain in terms of transnational access. The challenges of transnational access are many, starting from locating and documenting information on available sources and their content (i.e. coverage, variables, classifications, etc.) and on accreditation procedures (i.e. eligibility, rules, costs and timing). There are language barriers, as translated versions of information on data and accreditation procedures seldom exist or are incomplete. In addition, completing country-specific application forms for accreditation procedures is often demanding and different procedures exist for data held by different agencies even within the same country. Finally, data access systems differ across countries, implying that while remote access or execution could be possible in some countries, in others it is only possible to access on site, requiring researchers to travel to the location in question. These are just some of the challenges related to accessing data, before researchers can even begin confronting differences in the content and structure of micro-data themselves, and the time and human capital investment required to become acquainted with the "nitty gritty" of each database.

As a result, multi-country studies requiring the exploitation of micro-data are very difficult

Figure 1: Distributed Micro-data Analysis

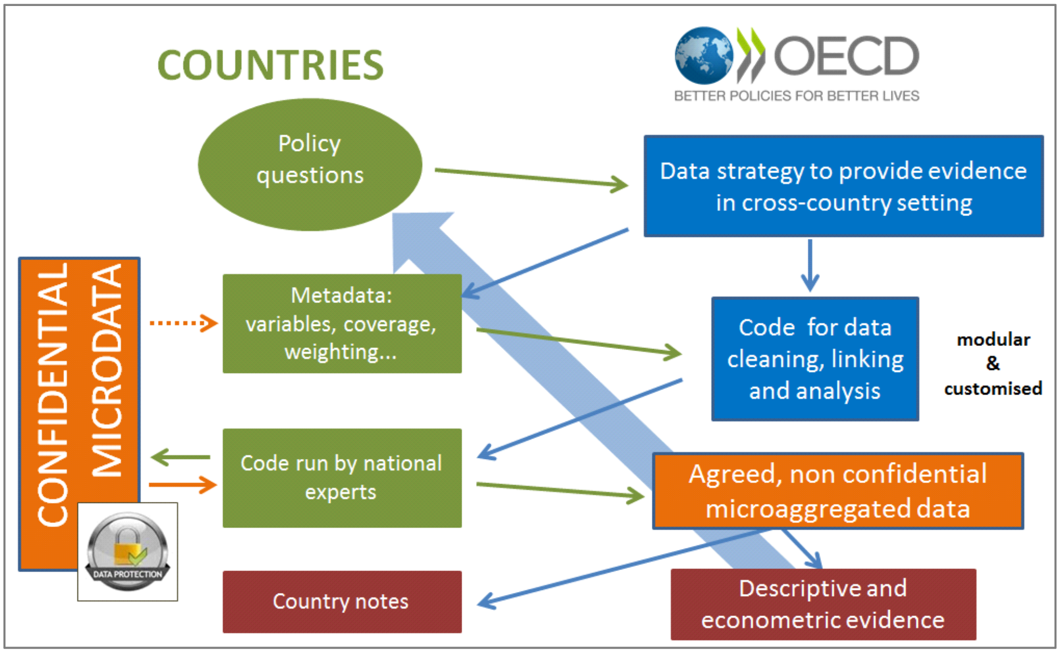

to conduct, and often rely on the formation and co-ordination of networks of national researchers, with each team having access to their respective national micro-data. The comparability of the country level results needs therefore to be insured via the use of a common protocol for data collection and aggregation and a common model specification for the econometric analysis.

The OECD pioneered this methodology, called distributed microdata analysis, at the beginning of the 2000s (OECD, 2003). It currently follows this approach in three ongoing projects: MultiProd, DynEmp, and MicroBeRD.[8] The distributed micro-data analysis involves running a common code in a decentralized manner by representatives in national statistical agencies or experts in public institutions, who have access to the national micro-level data. At this stage, micro-aggregated data are generated by the centrally designed, but locally executed, program codes, which are then sent for comparative cross-country analysis to the OECD. Figure 1 summarizes how the distributed micro-data approach works.

The advantages of this novel data collection methodology are manifold: it puts a lower burden on national statistical agencies and limits running costs for such endeavours. Importantly, it also overcomes the confidentiality constraints of directly using national micro-level statistical databases while at the same time achieving a high degree of harmonization and comparability across countries, sectors and over time.

In spite of these advantages, this procedure is still not widely applied today when collecting statistical information. This may have to do with the amount of time needed to set up and manage the network as well as to develop a well-functioning, "error-free" program code which is able

8 MultiProd, DynEmp, and MicroBeRD are projects carried out by the Directorate for Science, Technology and Innovation (STI) at the OECD. The DynEmp project provides harmonized microaggregated data to analyse employment dynamics (www.oecd.org/sti/dynemp) and MicroBERD provides information on R&D activity in firms from official business R&D surveys (www.oecd.org/sti/rdtax).

to both accommodate potential differences across national micro-level databases and minimize the burden on the researchers who have access to the data and run the code.

The MultiProd project is based on a distributed data collection exercise aimed at creating a harmonized cross-country micro-aggregated database on productivity patterns from confidential micro-level sources. In particular, the goal of the project is to investigate the extent to which different policy frameworks can shape firm productivity, and the way resources are allocated to more productive firms (i.e. allocative efficiency). Such analysis will be a key input for policy makers as firm-level productivity and efficient reallocation are the key engines of future growth.

The MultiProd Dataset

Variables and country coverage

The MultiProd program relies on two main data sources in each country. First, administrative data or production surveys (PS), which contain all the variables needed for the analysis of productivity but may be limited to a sample of firms. Second, a business register (BR), which contains a more limited set of variables but covers the entire population of firms. The program works also in the absence of a business register and this is indeed not needed when administrative data on the full population of firms are available. However, when data come from a PS, the representativeness of the results are substantially improved and, thus, their comparability across countries.

Census and administrative data, indeed, normally cover the whole population of businesses with at least one employee. Still, these datasets do not always exist and PS data need to be used.

Table 1: Temporal Coverage of the MultiProd Database, by Country

COUNTRY	YEAR																		
	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012
AUS									■	■	■	■	■	■	■	■	■	■	■
AUT															■	■	■	■	■
BEL										■	■	■	■	■	■	■	■	■	
CAN							■	■	■	■	■	■	■	■	■	■	■	■	■
CHL												■	■	■	■	■	■	■	■
DNK							■	■	■	■	■	■	■	■	■	■	■	■	■
FIN		■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
FRA		■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
HUN					■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
IDN									■	■	■	■	■	■	■	■	■	■	■
ITA								■	■	■	■	■	■	■	■	■	■	■	■
JPN	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	
LUX										■	■	■	■	■	■	■	■	■	■
NLD							■	■	■	■	■	■	■	■	■	■	■	■	■
NOR		■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■
NZL							■	■	■	■	■	■	■	■	■	■	■	■	
PRT											■	■	■	■	■	■	■	■	■
SWE									■	■	■	■	■	■	■	■	■	■	■

Source: MultiProd dataset, March 2017.

One of the big challenges of working with firm-level production surveys is that the selected sample of firms might yield a partial and biased picture of the economy. Whenever available, BRs, which typically contain the whole population of firms, are therefore used in MultiProd to compute a population structure by year-sector-size classes. This structure is then used to re-weight data contained in the PS in order to construct data that are as representative as possible of the whole population of firms and comparable across countries.

The MultiProd program computes a series of productivity measures that go from the least to the most data-demanding methodologies. To mention a few, gross output based LP; value added based LP; a Wooldridge (2009)-residual MFP based on value added as in Ackerberg *et al.* (2006); a Solow-residual based MFP using external, country-industry specific labour and intermediate shares; a Solow-residual based MFP using external, industry specific labour and intermediate shares (the cross-country-year median); a Superlative index based MFP using labour and intermediate shares calculated as the average between the labour/intermediate share of the firm (averaged over time) and the geometric mean of firm labour/intermediate shares in the industry.

For the MFP calculations a measure of capital stock is needed. In the baseline case, the program defines the capital stock variable through the perpetual inventory method (PIM) in order to increase the comparability of results across countries; the initial value is set to the capital stock reported by the firm in the initial year, whenever this is available. For countries that have capital stock information but not invest-

Table 2: Representativeness of the MultiProd Database for Manufacturing and Non-Financial Market Services, 2011

	Share of firms (%)		Share of employment (%)	
	BR	Eurostat	BR	Eurostat
Austria		69		92
Belgium		70		97
Denmark		102		117
Finland*	98	100	98	100
France		101		108
Hungary		94		99
Italy	11	11	54	54
Netherlands	6	4	56	44
Norway		71		89
Portugal**		92		97
Sweden		94		90

Source: MultiProd dataset (March 2017) and Eurostat, Business demography statistics

Note: Share of business registers and Eurostat data present in MultiProd. Manufacturing and non-financial market services only. The data come from different sources and the methodology used to treat the data (i.e., cleaning and calculation of sectors of activity) might differ; hence the comparison must be taken with caution. Shares higher than 100 per cent are likely to be due to these different methodologies adopted to treat data in the two sources. *: Finland is 100 per cent of firms with at least 1 FTE.

**: Data for Portugal include the population of active companies, but exclude individual enterprises (i.e. sole proprietors and self-employed); due to data limitations, we can only compare the MultiProd dataset to the total number of firms with at least one employee from Eurostat.

ment data, the capital stock becomes the main measure of capital. Finally, labour is measured by the number of employees or persons engaged (depending on data availability).

At the time of writing, 18 countries have been successfully included in the MultiProd database (namely, Australia, Austria, Belgium, Canada, Chile, Denmark, Finland, France, Hungary, Italy, Indonesia, Japan, Luxembourg, Netherlands, Norway, New Zealand, Portugal and Sweden). The data for each of the countries included so far are collected annually and at firm-level.

For most countries the time period covered by MultiProd spans between the early 2000s and 2012. Table 1 details for each country the exact period covered. MultiProd collects data for all sectors of the entire economy, whenever available. However, for the purposes of this analysis we have restricted our sample to the manufacturing and non-financial market services sector.

To provide an idea of the coverage for the European countries contained in the MultiProd dataset, Table 2 reports for 2011 the share of firms and employment with respect to both the Business Register (when available) and to Eurostat data (annual business demography by size class database). The table is constructed for the manufacturing and non-financial market service sectors. The data from Eurostat refer to the total number of firms or the total number of firms with at least one employee, in accordance with the micro-data used in MultiProd.

Comparing across different data sources is never easy, but data from Eurostat give a good benchmark to compare our data. The coverage is rather high in most of the countries (and results are very similar for each year of the sample). In particular, we have datasets covering roughly the population of firms for all countries reported in the table, except for Italy and the Netherlands. However, for these two countries the full BR is available, and thus the samples are reweighted. For instance, Italy has a skewed distribution with a large mass of very small firms which cannot be captured by production surveys. The survey used by MultiProd contains only 11 per cent of the total population of firms (both with respect to the BR and to the data published by Eurostat) but it accounts for 54 per cent of total employment. At the same time we have access to the entire population of firms from the Business Register, which we use to reweight our sample moments.

In the Netherlands the situation is similar, with the only existing survey of firms representing a very small share of firms, but the BR allows us to re-weight those firms ex-post in order to make the reported statistics representative of the total economy. In other words, in all countries except Italy and the Netherlands each firm has a weight equal (or close) to one, whereas the Italian and the Dutch datasets have been reweighted using the BR, which cover the population of firms.[9]

The weighting procedure entails the following two main steps:

1) Preparation of the population structure from the Business Registry (BR): the number of firms by year, industry, size class is obtained from the BR, using the most detailed industry level available and seven size classes (with thresholds at 5, 10, 20, 50, 100 and 250 employees).

2) Calculation of actual weights: the weights are computed, for each variable, as the number of firms in the population of the corresponding industry and size class divided by the number of firms in the survey, after having cleaned the data through an outlier filter.

9 The weights are variable-specific, hence missing information or outliers might cause weight to be different from one even in the presence of data containing the entire population of businesses.

Output

MultiProd output is a collection of files that contain statistics for different variables computed yearly at a detailed sectoral level.[10] The program is flexible to allow for multiple levels of aggregation at which the output is produced and, for each of those levels, the types of aggregated data to be included.

The statistics are never collected at the level of the individual firm. Instead the programme splits firms along various dimensions, into cells, and for each cell collects aggregate annual data. In addition, these statistics are collected in terms of both levels and growth rates. The dimensions used for the split are the following, which can be specified at different levels of aggregation:

- Sector: 1-digit (STAN A7) aggregation level and 2-digit (STAN A38) aggregation level;
- Firm-level productivity distribution: splitting firms into productivity quantiles (with productivity defined in various ways, such as LP, MFP à la Wooldridge 2009, MFP à la Solow);
- Gross output distribution: splitting firms into quantiles based on gross output;
- Size class: splitting firms into groups based on employment levels;
- Age of the firm: splitting firms into groups based on age;
- Ownership: independent firm vs. affiliate of a business group, and nationality of the group;
- Demographics: entrants, exitors, incumbents, etc.

The output is provided also combining the previous dimensions together (e.g. the interaction of age and size classes).[11]

Several statistics are collected:

- Basic moments: mean, median, standard deviation, and number of non-missing values, for a series of variables.
- Several measures of aggregate productivity: decomposition of both aggregate LP and/or MFP, together with allocative efficiency measures (Olley-Pakes 1996 covariance terms).
- Measures of allocative efficiency based on Hsieh and Klenow (2009), and Petrin and Sivadasan (2013) to analyse the role of allocation and selection.
- Distribution characteristics for productivity levels, productivity growth rates, output and wages. These include both non-parametric measures such as percentiles and parameters of the distributions (e.g., pareto).
- Descriptive statistics of firm characteristics (including growth rate and wage dispersion) by quantiles of the productivity distribution in levels and growth, and by quantiles of the sales distribution.
- Characteristics (productivity, age, persistence, size, ownership, investments etc.) of firms at the productivity frontier.
- Employment dynamics by quantiles of the productivity distribution.
- Estimated parameters from distributed regressions at the firm level, with the aim of establishing a set of stylised facts for each country regarding the relationship between productivity, firm characteristics (size, age, previous performance, ownership, etc.) and structural characteristics (concentration, misallocation, sectoral policies, etc.).
- Tabulations of firms with negative value added and graphs of the sectors' productivity distributions.

The output produced by the program covers a wide range of topics: productivity heterogene-

10 For further details see Berlingieri *et al.*(2017b).

11 However note that, although possible, the code never combines more than three dimensions at the same time; the reason is that the number of firms in each cell would become small enough to incur confidentiality problems, especially at high levels of industry disaggregation.

ity; allocative efficiency; granularity; and wage dispersion and its link to productivity.

In order to examine the effects of policies and macro shocks on productivity heterogeneity we collect information about the distribution of productivity (using different measures for it, as described before). In particular, we aggregate productivity and its variance to the sector level, and we decompose productivity dispersion into within- versus between-sector and within- versus between- quantile dispersion. Given the relevance of this topic, this article is focused exclusively on this block of results.

The access to firm-level data across multiple countries also allows us to conduct an extensive examination of allocative efficiency over time and across countries, applying a number of different methods. The methods used are: the Olley and Pakes (1996) static productivity decomposition, as well as a dynamic version of it (Melitz and Polanec, 2015); measurement of job reallocation; measurement of productivity dispersion in the top and bottom size quantile of firms in each sector and comparison to the productivity dispersion in the whole sector; measurement of misallocation according to the Hsieh and Klenow (2009) procedure with some refinements; a description of the distribution of firm-level distortions in each input and their overall impact on aggregate productivity; analysis of the gap between the value of the marginal product of an input and its marginal cost as in Petrin and Sivadasan (2013); and run of distributed firm-level regressions of these measures of misallocation on firm characteristics such as size, age, and ownership.

In addition, the role of the largest firms can be investigated in great detail and compared to that of the most productive firms ("frontier firms"). This analysis sheds light on the so-called "granular" hypothesis, which posits that aggregate fluctuations are the results of microeconomic shocks and not economy-wide shocks, as usually assumed.[12] Such idiosyncratic shocks, even if they are uncorrelated, may not cancel on average if sectors are dominated by a small number of large firms. This fact can have important implications for policies aimed at increasing economic resilience, and highlights the importance of studying firm-level data to better understand aggregate outcomes. The MultiProd project can offer new insights on this hypothesis by analysing how much of a country's economic activity is driven by a small number of important firms, and how much of the observed productivity variation is indeed the result of microeconomic variations. The program collects a number of indicators, such as: the market share and the share of employment accounted for by the top decile of firms in terms of gross output and productivity; the sectoral level Hirsch-Herfindahl Index (HHI); the decomposition of aggregate productivity in both level and growth between the contribution of the largest firms and that of the other firms.

Finally, data collected in the MultiProd project are instrumental to understand the evolution of the between-firm component of wage dispersion, which has been found to account for a large share of the wage inequality of individuals (see, for example, Dunne *et al*., 2004, Card *et al*., 2013, Card *et al*., 2014, Song *et al*., 2015). In particular the program decomposes the wage dispersion in the within and between contribution both for industry and productivity quantiles; calculates the share of each industry and productivity quantile in the overall within component of wage dispersion; identifies the impact

12 With "granularity" we refer to the extent to which economic activity in general, and aggregate productivity in particular, is driven by a small number of large firms. When large firms represent a disproportionate share of the economy, indeed, aggregate fluctuations may be governed by idiosyncratic shocks to these large firms. This hypothesis - called the "granular hypothesis" and proposed in Gabaix (2011) - suggests that aggregate fluctuations can be traced back to micro shocks hitting a small number of large firms.

of various policies (e.g. minimum wage) on dispersion through distributed regressions.

The next section will focus on productivity heterogeneity, and provide evidence on the dispersion of productivity and its evolution over time obtained with the MultiProd dataset.

Productivity Heterogeneity

A startling fact of firm-level productivity analysis is the large and persistent differences in both LP and MFP between firms, even within narrowly defined sectors.

The MultiProd project offers a detailed understanding of productivity dispersion by investigating the relationship between microeconomic dispersion and economy-wide dispersion, in order to provide a better illustration of productivity variation within countries. In particular, to document productivity heterogeneity, MultiProd collects information about the distribution of productivity, aggregates productivity and its variance at the 2-digit sector and at the macro-sectoral level, and decomposes productivity dispersion into within — versus between — sector components.

Productivity Dispersion and its Evolution over Time

In order to capture heterogeneity in the data, MultiProd calculates several measures of dispersion for productivity within macro-sectors and 2-digit industries: the standard deviation; the 90-10, 90-50, and 50-10 ratios; as well as the interquartile range (i.e., the difference between the 75th and the 25th percentile).

In particular, the 90-10 productivity ratio is defined as the ratio of the 90th percentile to the 10th percentile of the productivity distribution. It is used widely in the productivity literature to assess the spread of the productivity distribution. The measure is quite intuitive since a 90-10 ratio of X can be interpreted as firms at the top of the productivity distribution, proxied by firms at the 90th percentile, producing X times as much as firms at the bottom of the distribution, proxied by firms at the 10th percentile, given the same amount of inputs.

As an example, in Table 3 we illustrate the 90-10 ratio for both (log) LP and (log) MFP in 2011. The table illustrates some important features. First, there is a rather significant produc-

Table 3: Productivity 90-10 Ratio in 2011, by Country

	Log(LP_VA) 90-10 ratio		Log(MFP) 90-10 ratio	
	Manufacturing	Services	Manufacturing	Services
Australia	1.87	2.05	1.90	2.12
Austria	1.96	2.42		
Belgium	1.60	1.74	1.80	1.78
Chile	3.00	3.53	3.07	3.87
Denmark	1.46	1.96	1.32	1.80
Finland	1.17	1.38	1.19	1.34
France	1.35	1.81	1.40	1.78
Hungary	2.79	3.29	2.54	2.86
Indonesia	3.11	-	3.41	-
Italy	1.66	2.01	1.60	1.88
Japan	1.26	1.38	1.17	1.38
Netherlands	2.00	2.98	2.27	2.27
New Zealand	1.84	2.09	1.92	2.08
Norway	1.73	2.17	1.87	2.08
Portugal	1.88	2.65	1.88	2.75
Sweden	1.45	1.86	1.59	2.45
Unweighted Average	1.88	2.22	1.93	2.17

Source: MultiProd dataset, March 2017.

Chart 1: Top, Median and Bottom Decile Over Time

Panel A: (log) LP

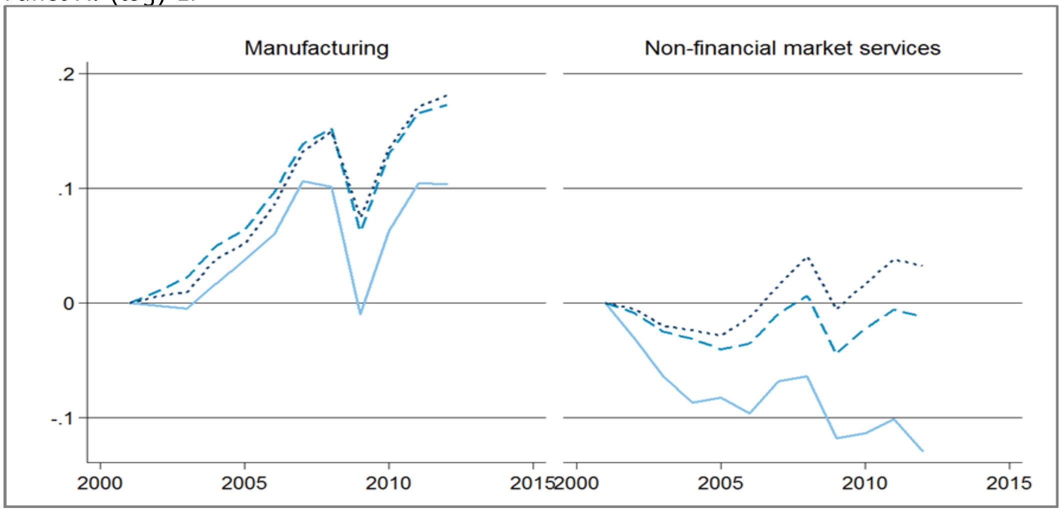

Panel B: (log) MFP

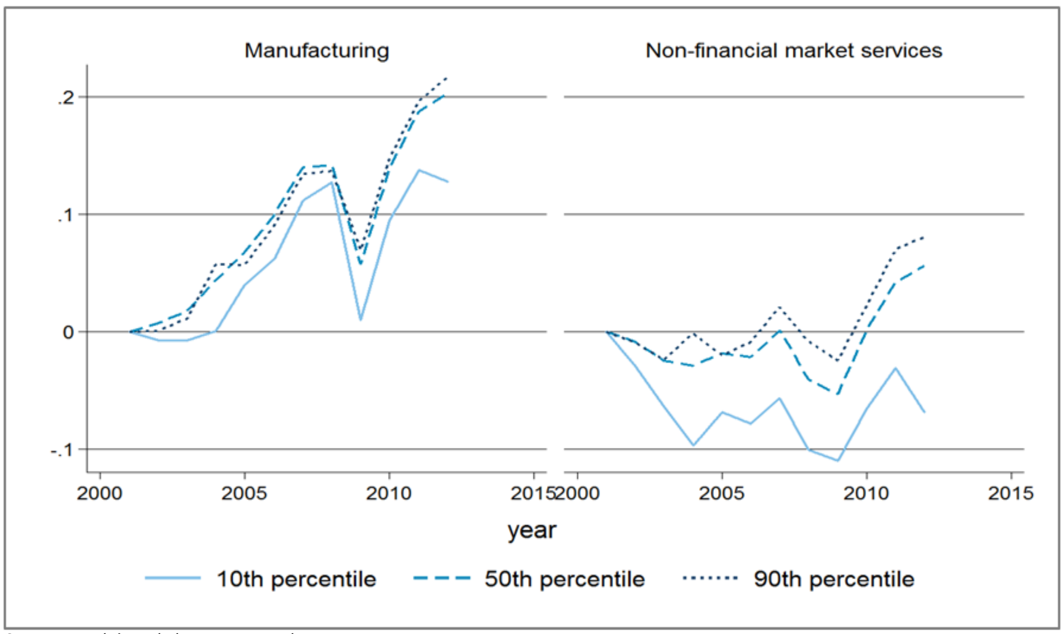

Source: MultiProd dataset, March 2017.

Note: Lop-LP (panel A) and Log-MFP à la Wooldridge (panel B) in the 10th, 50th and 90th percentile of the productivity distribution, for manufacturing (left panel) and services (right panel) since 2000. The countries included are AUS, AUT, BEL, CHL, DNK, FIN, FRA, HUN, ITA, JPN, NLD, NOR, NZL, PRT, SWE. IDN is available only for manufacturing; therefore, for comparability across sectors, it has been excluded from the graph. The graphs can be interpreted as the cumulated growth rates of LP (MFP) within each country and sector over the period. For instance, in 2012 in manufacturing in the 90th quantile of productivity is roughly 19% (24%) higher than in 2001. The estimates reported in Panel A (Panel B) are those of year dummies in a cross-country regression of log-LP (log-MFP) in the 90th, 50th and 10th percentile of the distribution.

tivity dispersion in both manufacturing and services between the top performing and the bottom performing firms, and both in terms of LP and MFP. Second, dispersion is on average higher in services than in manufacturing, whether in terms of LP or MFP. Third, the ratio is particularly high in Chile, Indonesia and Hungary. Finally, in 2011, on average across countries, firms in the top decile of the distribution can produce more than six times as much value added per worker as firms in the bottom decile of the same country's manufacturing sector, and nine times in services. Essentially the same proportion is kept when looking at the 90-10 ratio in terms of MFP: with the same amount of measured inputs, firms at the top of the distribution produce almost seven times the output of firms at the bottom in the manufacturing sector, and almost nine times in the services sector. The large dispersion reflects the heterogeneity of the sample, with developing countries such as Indonesia displaying a very high dispersion.[13]

Chart 1 shows the productivity (LP and MFP, respectively) of the 10th, the 50th and 90th percentile of the (log) productivity distribution, normalizing the year 2001 to 0. In each figure, the left panel represents productivity dispersion in manufacturing and the right panel represents productivity dispersion in (non-financial) market services. The data show that there has been an increase in dispersion of productivity over time, especially in the services sector. The negative effect that the Great Recession had on the productivity trends is also evident from the figures, especially at the bottom of the distribution.

Productivity Dispersion Decomposition

To better understand the origin of the productivity dispersion, it is possible to decompose the total econmy productivity variance V_t at a point in time into two components: a within-industry component V_{Ft} and a cross-industry component V_{Xt}. The within-industry component V_{Ft} captures how much a firm's individual productivity differs from the sector (labour-weighted) average. The cross-industry component V_{Xt} captures instead how much sectors vary from each other.

$$V_t = V_{Ft} + V_{Xt} \quad (1)$$

The within-industry variance V_{Ft} is the average over all sectors j of the square deviation of firms' productivity P_{it} to their sector (weighted) average LP $\bar{P}_{jt}$:

$$\begin{aligned} V_{Ft} &= \frac{1}{L_t}\sum_j \sum_{i \in j} L_{it}(P_{it} - \bar{P}_{jt})^2 \\ &= \sum_j \frac{L_{jt}}{L_t} \sum_{i \in j} \frac{L_{it}}{L_{jt}}(P_{it} - \bar{P}_{jt})^2 \\ &= \sum_j \frac{L_{jt}}{L_t} \delta_{jt}^2 \end{aligned} \quad (2)$$

and the cross-industry component V_{Xt} is the average of the squared deviation of sector j's average productivity $\bar{P}_{jt}$ to the economy-wide productivity $\bar{P}_t$:

$$V_{Xt} \equiv \sum_j \frac{L_{jt}}{L_t}(\bar{P}_{jt} - \bar{P}_t)^2 \quad (3)$$

where L_{jt}/L_t is the employment share of sector j at time t, and $\delta_{jt}^2 \equiv \sum_{i \in j} \frac{L_{it}}{L_{jt}}(P_{it} - \bar{P}_{jt})^2$ is the labour-weighted industry variance of firm-level LP.[14]

The MultiProd project contributes to the literature by offering a detailed decomposition of overall productivity dispersion based on cross-country microeconomic data. This decomposition can help understand how much of the country-level dispersion in productivity comes from

13 The table display the average dispersion within 2-digit sectors. While performing the same exercise within 4-digit sectors would partially reduce the dispersion (not available due to confidentiality), the comparison across countries at the 2-digit is nevertheless informative. Moreover, the microdata used for MultiProd contain, or are representative for, the whole population of firms with at least one employee, which naturally imply a higher level of dispersion than other more selected samples.

Table 4: Share of Within-Sector LP Variance in Total Productivity Dispersion, 2011

	LP Dispersion	
	Manufacturing	Services
Australia	98	99
Austria	86	90
Belgium	76	88
Chile	90	97
Denmark	84	63
Finland	65	76
France	63	85
Hungary	79	99
Indonesia	79	-
Italy	82	65
Japan	75	89
Netherlands	80	71
Norway	83	73
Portugal	62	76
Sweden	53	74
Unweighted Average	77	82

Source: MultiProd dataset, March 2017.

microeconomic dispersion within narrowly defined sectors, and how much comes from more aggregate shocks that affect whole sectors. This is achieved by looking at the ratio V_{Ft}/V_t which reflects the importance of microeconomic shocks to aggregate dispersion. The decomposition suggested here is a cross-sectional decomposition of productivity dispersion in a given period t.

Table 4 presents the ratio V_{Ft}/V_t for LP in 2011. The two columns report the share of total LP dispersion accounted for by within-sector dispersion, for manufacturing and services respectively. The results show that on average within-sector dispersion accounts for more than 77 per cent (82 per cent) of the overall LP dispersion observed across firms in manufacturing (services): a large share of dispersion comes from heterogeneity in LP between firms within the same two-digit sector. In other words, a substantial part of productivity heterogeneity does not come from the type of activity that firms engage in, per se, but rather from more intrinsic differences among firms within the same sector of activity in the same country.[15]

In addition to the above decomposition, the overall within-industry component can then be further decomposed into the contribution from each industry, to precisely pin down which industry drives productivity dispersion, i.e. which are the industries where dispersion is stronger. Similarly to Carvalho and Gabaix

14 Note that this is the variance decomposition of LP. It can be generalized to MFP but the choice of the appropriate weights becomes less straightforward. In the literature it is common to use output weights (gross output or value added, depending on how MFP is estimated) but the resulting weighted average does not correspond to the precise measure of aggregate productivity. Moreover, the standard Domar weights used to decompose (gross output) MFP productivity growth do not yield an exact decomposition. Van Biesebroeck (2008) shows that to do so one would need more complex input weights.

15 Table 4 displays the average dispersion within 2-digit sectors. As already stated in footnote 11 for the previous exercise, performing this decomposition within 4-digit sectors would partially reduce the share of within-sector productivity variance (not available to us due to confidentiality). The comparison across countries at the 2-digit level is nevertheless informative. Moreover, the microdata used for MultiProd contain, or are representative for, the whole population of firms with at least one employee, which naturally imply a higher level of dispersion than other more selected samples.

Table 5: Contribution of the Top Three Sectors in the Share of Within-Sector LP Variance in Total Productivity Dispersion, 2011

	Manufacturing	Var. Share	Services	Var. Share
Australia	Food products, beverages and tobacco [CA]	35	Transportation and storage [H]	26
	Machinery and equipment n.e.c. [CK]	28	Legal and accounting activities, etc [MA]	25
	Basic pharmaceutical products and pharmaceutical preparations [CF]	13	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	21
Austria	Machinery and equipment n.e.c. [CK]	17	REAL ESTATE ACTIVITIES [L]	34
	Basic metals and fabricated metal products, except machinery and equipment [CH]	15	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	25
	Wood and paper products, and printing [CC]	12	Administrative and support service activities [N]	1
Belgium	Chemicals and chemical products [CE]	42	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	35
	Food products, beverages and tobacco [CA]	30	Legal and accounting activities, etc [MA]	22
	Wood and paper products, and printing [CC]	7	Transportation and storage [H]	22
Chile	Wood and paper products, and printing [CC]	54	REAL ESTATE ACTIVITIES [L]	39
	Food products, beverages and tobacco [CA]	22	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	36
	Coke and refined petroleum products [CD]	11	Transportation and storage [H]	10
Denmark	Machinery and equipment n.e.c. [CK]	28	Transportation and storage [H]	25
	Furniture; other manufacturing; repair and installation of machinery and equipment [CM]	24	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	22
	Food products, beverages and tobacco [CA]	15	REAL ESTATE ACTIVITIES [L]	13
Finland	Computer, electronic and optical products [CI]	28	REAL ESTATE ACTIVITIES [L]	64
	Wood and paper products, and printing [CC]	17	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	18
	Basic pharmaceutical products and pharmaceutical preparations [CF]	12	Telecommunications [JB]	5
France	Basic pharmaceutical products and pharmaceutical preparations [CF]	17	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	22
	Computer, electronic and optical products [CI]	17	Legal and accounting activities, etc [MA]	22
	Food products, beverages and tobacco [CA]	12	Transportation and storage [H]	11
Hungary	Transport equipment [CL]	32	Administrative and support service activities [N]	36
	Computer, electronic and optical products [CI]	25	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	23
	Machinery and equipment n.e.c. [CK]	23	REAL ESTATE ACTIVITIES [L]	15
Indonesia	Transport equipment [CL]	34		
	Chemicals and chemical products [CE]	18		
	Food products, beverages and tobacco [CA]	17		
Italy	Machinery and equipment n.e.c. [CK]	15	Telecommunications [JB]	23
	Transport equipment [CL]	13	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	22
	Basic metals and fabricated metal products, except machinery and equipment [CH]	12	Transportation and storage [H]	14
Japan	Coke and refined petroleum products [CD]	25	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	37
	Chemicals and chemical products [CE]	17	Legal and accounting activities, etc [MA]	34
	Machinery and equipment n.e.c. [CK]	11	Administrative and support service activities [N]	11
Netherlands	Chemicals and chemical products [CE]	34	Administrative and support service activities [N]	35
	Food products, beverages and tobacco [CA]	10	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	28
	Coke and refined petroleum products [CD]	10	Telecommunications [JB]	12

Norway	Machinery and equipment n.e.c. [CK]	28	Telecommunications [JB]	31
	Electrical equipment [CJ]	14	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	22
	Transport equipment [CL]	11	Transportation and storage [H]	16
Portugal	Coke and refined petroleum products [CD]	28	REAL ESTATE ACTIVITIES [L]	21
	Food products, beverages and tobacco [CA]	20	Telecommunications [JB]	21
	Wood and paper products, and printing [CC]	17	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	18
Sweden	Computer, electronic and optical products [CI]	35	REAL ESTATE ACTIVITIES [L]	39
	Wood and paper products, and printing [CC]	14	Telecommunications [JB]	16
	Transport equipment [CL]	12	Wholesale and retail trade, repair of motor vehicles and motorcycles [G]	15

Source: MultiProd dataset, March 2017.

(2013), who investigate the importance of granularity - or microeconomic shocks - in driving macroeconomic fluctuations, one can decompose the within-industry component of the productivity variance into the weighted productivity variances of industries, as shown in the last term of Equation (2).[16]

We report in Table 5 the top three contributors to LP variance in 2011 for each country, and for manufacturing and services respectively. Some sectors, such as "food products, beverage and tobacco", "machinery and equipment", "wholesale and retail trade, repair of motor vehicles and motorcycles," and "transport and storage" regularly appear amongst the sectors characterized by the highest productivity dispersion. This suggests that there might be sectoral features of the within-sector distribution of firms that might affect the distribution of productivity.

We now describe how the share of within-sector variance of LP evolves over time, particularly in light of the Great Recession. The results displayed in Chart 2 suggest that within sector variance of LP remained the most important component of overall variance, well above 75 per cent, but its importance declined in both manufacturing and services after 2008. In other words, this suggests that in the aftermath of the crisis a larger share of the productivity dispersion came from productivity differences across rather than within sectors. This might suggest that the aggregate shock of the Great Recession might have affected more systematically certain sectors, such as durables, relative to how systematically it affected firms at the top and the bottom of the productivity distribution within sectors. Nonetheless, this impact still left a large part of productivity heterogeneity that cannot be explained by sectoral differences, suggesting that cross-sectoral analyses are likely to underestimate the amount of productivity divergence in the economy.

Conclusions and Avenues for Future Research

This article provides an overview of the main contributions of the MultiProd project in light of the current literature. It focuses, in particular, on the role of productivity heterogeneity and the evidence that emerges from the distributed micro-data analysis carried out in the project, which resulted in a unique dataset of harmonized cross-country moments that are representative for the population of firms and

16 This is an exact decomposition of the within-industry component of variance, which differs from what they define as fundamental volatility for the weights (not squared in the present case) and the variance (computed on the cross section of firms and not constant over time).

Chart 2: Share of Within-Sector log-LP Dispersion in Total Productivity Dispersion, 2000-2012

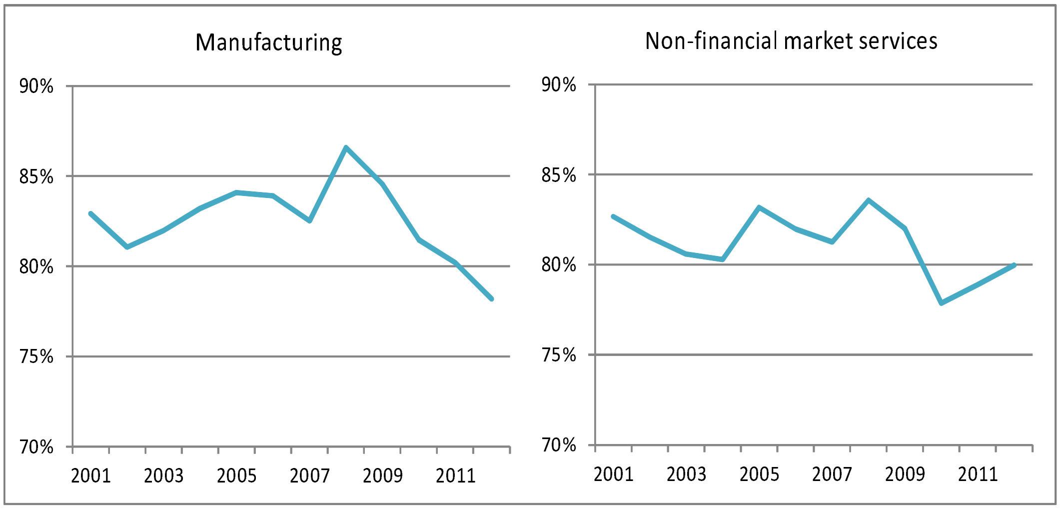

Source: MultiProd dataset, March 2017.

Note: Share of within-sector dispersion in overall macro-sector Log-LP dispersion. Average across countries and sectors, weighted by employment. Countries: AUS, AUT, BEL, CHL, DNK, FIN, FRA, HUN, ITA, JPN, NLD, NOR, NZL, PRT, SWE.

comparable across countries even at a detailed two-digit industry level.

We have shown that productivity dispersion is high in both manufacturing and non-financial market services: in 2011, for instance, firms in the top decile of the distribution produced more than six times as much value added per worker as firms in the bottom decile of the same country's manufacturing sector, and nine times in services. Moreover, dispersion has increased over time, especially in services. A substantial part of this productivity heterogeneity seems to come from differences among firms within the same sector of activity in the same country, rather than from the type of activity that firms engage in, per se. On average, the within-sector dispersion accounts for more than 77 per cent (82 per cent) of the overall LP dispersion observed across firm in manufacturing (services). Finally, the analysis of the within sector variance of LP over time suggests that it remained the most important component of the overall variance, well above 80 per cent, for almost the entire period. From 2008 onward its importance slightly declined both in manufacturing and services, suggesting that the Great Recession might have systematically affected certain sectors, more than systematically affecting firms at the top and the bottom of the productivity distribution within sectors. In any case, within-sector variance remains by far the main component of the overall LP variance even after 2008.

This article has focused on productivity heterogeneity, but, thanks to the richness of the output of the MultiProd project, other analyses can be carried out in order to:

- Better depict the Schumpeterian process of creative destruction across countries;
- Gauge whether resources are efficiently allocated through the analysis of the firm-level productivity distribution, with further refinements by size, age, and ownership categories;
- Identify the largest firms and understand how they differ in terms of their weight in the economy, their productivity performance, and their contribution to aggregate productivity growth;

- Identify firms at the 'frontier' - the best performers - and understand how they differ across countries, what drives their performance, and how much they contribute to aggregate productivity growth;
- Investigate the cross-country differences in firm-level productivity performance and allocative efficiency before, during and after the financial crisis;
- Investigate the relationship between productivity and wage dispersion, and gauge to what extent heterogeneity in productivity has contributed to wage inequality; and
- Examine the effectiveness of various policy frameworks aimed at shaping firm productivity and enhancing resource allocation to more productive firms.

These are just some of the possible interesting avenues that we plan to address in subsequent work. Last but not least we aim at linking differences in these important features of productivity dynamics and distributions to structural changes, such as digitalization and globalization, and country framework conditions and policies.

References

Ackerberg, Daniel, Kevin Caves, and Garth Frazer (2006) "Structural Identification of Production Functions," http://mpra.ub.uni-muenchen.de/38349/.

Arnold, Jens, Giuseppe Nicoletti, and Stefano Scarpetta (2008) "Regulation, Allocative Efficiency and Productivity in OECD Countries: Industry and Firm-Level Evidence," OECD Economics Department Working Papers, No. 616.

Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta (2009) "Measuring and Analyzing Cross-Country Differences in Firm Dynamics," in *Producer Dynamics: New Evidence from Micro Data*, pp. 15-76. (Chicago: University of Chicago Press).

Bartelsman, Eric, Stefano Scarpetta, and Fabiano Schivardi (2005) "Comparative Analysis of Firm Demographics and Survival: Evidence from Micro-Level Sources in OECD Countries," OECD Economics Department Working Papers, No. 348.

Berlingieri, Giuseppe, Patrick Blanchenay, and Chiara Criscuolo (2017a) "The Great Divergence(s)," OECD Science, Technology and Industry Policy Papers, No. 39.

Berlingieri, Giuseppe, Patrick Blanchenay, Sara Calligaris and Chiara Criscuolo (2017b) "The MultiProd project: A Comprehensive Overview," OECD Science, Technology and Industry Working Papers, No. 2017/04.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts (2013) "Does Management Matter? Evidence from India," *Quarterly Journal of Economics*, Vol. 128, No. 1, pp. 1-51.

Breschi, Stefano, Franco Malerba, and Luigi Orsenigo (2000) "Technological Regimes and Schumpeterian Patterns of Innovation," *Economic Journal*, Vol. 110, No. 463, pp. 388-410.

Card, David, Francesco Devicienti, and Agata Maida (2014) "Rent-Sharing, Holdup, and Wages: Evidence from Matched Panel Data," *Review of Economic Studies*, Vol. 81, No.1, pp. 84-111.

Card, David, Jörg Heining, and Patrick Kline (2013) "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics*, Vol. 128, No. 3, pp. 967-1015.

Carvalho, Vasco M., and Xavier Gabaix (2013) "The Great Diversification and Its Undoing," A*merican Economic Review*, *Vol.* 103, No. 5, pp. 1697-1727.

Caselli, Francesco (1999) "Technological Revolutions," *American Economic Review*, Vol. 89, No. 1, pp. 78-102.

Comin, Diego, and Bart Hobijn (2006) An Exploration of Technology Diffusion. National Bureau of Economic Research Working Paper No. 12314.

Dunne, Timothy, Lucia Foster, John Haltiwanger, and Kenneth R. Troske (2004) "Wage and Productivity Dispersion in United States Manufacturing: The Role of Computer Investment," J*ournal of Labor Economics*, Vol. 22, Nol. 2, pp. 397-429.

Foster, Lucia, Cheryl Grim, and John Haltiwanger (2014) "Reallocation in the Great Recession: Cleansing or Not?" National Bureau of Economic Research. Working Paper No. 20427.

Foster, Lucia, John C. Haltiwanger, and Cornell John Krizan (2001) "Aggregate Productivity Growth. Lessons from Microeconomic Evidence," in *New Developments in Productivity Analysis*, pp. 303-72. (Chicago: University of Chicago Press).

Foster, Lucia, John Haltiwanger, and C. J. Krizan (2002) "The Link Between Aggregate and Micro Productivity Growth: Evidence from Retail

Trade," National Bureau of Economic Research, Working Paper No. 9120.

Gabaix, Xavier (2011), "The Granular Origins of Aggregate Fluctuations", Econometrica, Vol. 79, No. 3, pp. 733-772.

Hall, Robert E., and Charles I. Jones (1999) "Why Do Some Countries Produce So Much More Output Per Worker than Others?" *Quarterly Journal of Economics*, Vol. 114, No. 1, pp. 83-116.

Hopenhayn, Hugo A (2014) "Firms, Misallocation, and Aggregate Productivity: A Review," *Annual Review of Economics*, Vol. 6, No. 1, pp. 735-770.

Hsieh, Chang-Tai, and Peter J. Klenow (2009) "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, Vol. 124, No. 4, pp. 1403-48.

Hsieh, Chang-Tai, and Peter J. Klenow (2010) "Development Accounting," *American Economic Journal: Macroeconomics* Vol. 2 No. 1, pp. 207-23.

Klenow, Peter and Andres Rodriguez-Clare (1997) "The Neoclassical Revival in Growth Economics: Has It Gone Too Far?" in *NBER Macroeconomics Annual*, Vol. 12, pp. 73-114. (Cambridge: MIT Press).

Malerba, Franco and Luigi Orsenigo (1995) "Schumpeterian Patterns of Innovation," *Cambridge Journal of Economics*, Vol. 19, No. 1, pp. 47-65.

Malerba, Franco and Luigi Orsenigo (1996) "Schumpeterian Patterns of Innovation are Technology-Specific," *Research Policy*, Vol. 25, No. 3, pp. 451-478.

Melitz, Marc J. and Sašo Polanec (2015) "Dynamic Olley-Pakes Productivity Decomposition with Entry and Exit," *RAND Journal of Economics*, Vol. 46, No. 2, pp. 362-375.

OECD (2003) *The Sources of Economic Growth in OECD Countries*, (Paris: OECD Publishing).

Olley, G. Steven and Ariel Pakes (1996) "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, Vol. 64, No. 6, pp. 1263-97.

Parente, Stephen L. and Edward C. Prescott (1994) "Barriers to Technology Adoption and Development," *Journal of Political Economy*, Vol. 102, No. 2, pp. 298-321.

Petrin, Amil and Jagadeesh Sivadasan (2013) "Estimating Lost Output From Allocative Inefficiency, with an Application to Chile and Firing Costs," *Review of Economics and Statistics*, Vol. 95, No. 1, pp. 286-301.

Prescott, Edward C (1998) "Lawrence R. Klein Lecture 1997: Needed: A Theory of Total Factor Productivity," *International Economic Review*, Vol. 11, No. 4, pp. 525-51.

Restuccia, Diego, and Richard Rogerson (2013) "Misallocation and Productivity," *Review of Economic Dynamics*, Vol. 16, No. 1, pp. 1-10.

Schmitz Jr, James A (2005) "What Determines Productivity? Lessons from the Dramatic Recovery of the US and Canadian Iron Ore Industries Following Their Early 1980s Crisis," *Journal of Political Economy*, Vol. 113, No. 3, pp. 582-625.

Schwellnus, Cyrille and Jens Arnold (2008) "Do Corporate Taxes Reduce Productivity and Investment at the Firm Level?: Cross-Country Evidence from the Amadeus Dataset," OECD Economics Department Working Papers, No. 641, (Paris: OECD Publishing)

Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter (2015) "Firming Up Inequality," Working Papers 21199, National Bureau of Economic Research.

Syverson, Chad (2004) "Product Substitutability and Productivity Dispersion," *Review of Economics and Statistics*, Vol. 86, No. 2, pp. 534-50.

Van Biesebroeck, Johannes (2008) "Aggregating and Decomposing Productivity," *Review of Business and Economics*, Vol. 53, No. 2, pp. 122-46.

van Dijk, Machiel (2000) "Technological regimes and industrial dynamics: the evidence from Dutch manufacturing," *Industrial and Corporate Change*, Vol. 9, No. 2, pp. 173-194.

Wooldridge, Jeffrey M (2009) "On Estimating Firm-Level Production Functions Using Proxy Variables to Control for Unobservables," *Economics Letters*, Vol. 104, No. 3, pp. 112-14.