

NBER WORKING PAPER SERIES

NEW EVIDENCE ON SECTORAL LABOR PRODUCTIVITY:
IMPLICATIONS FOR INDUSTRIALIZATION AND DEVELOPMENT

Berthold Herrendorf
Richard Rogerson
Ákos Valentinyi

Working Paper 29834
<http://www.nber.org/papers/w29834>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2022

We have received helpful comments and suggestions from Gaaitzen de Vries, Douglas Gollin, and the participants of presentations at ASU, CEPR's Annual STEG Conference 2022, and Swansea University. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w29834.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Berthold Herrendorf, Richard Rogerson, and Ákos Valentinyi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

New Evidence on Sectoral Labor Productivity: Implications for Industrialization and Development
Berthold Herrendorf, Richard Rogerson, and Ákos Valentinyi
NBER Working Paper No. 29834
March 2022
JEL No. E24,O11,O14,O47

ABSTRACT

Moving labor from agriculture to manufacturing – “industrialization” – is often viewed as essential for the development of poor countries. We present new evidence on the channels through which industrialization can help poor countries close the productivity gap with rich countries. To achieve this, we leverage recent data releases by the Groningen Growth and Development Centre and build a new dataset of comparable labor productivity levels in agriculture and manufacturing for 64 mostly poor countries during 1990–2018. We find two key results: (i) cross-country labor productivity gaps in manufacturing are larger than in the aggregate and (ii) there is no tendency for manufacturing labor productivity to converge. While these results challenge the notion that expanding manufacturing employment is essential for the development of today’s poor countries, we also find that higher labor productivity growth in manufacturing is associated with higher labor productivity growth in the aggregate and in several key sectors.

Berthold Herrendorf
W.P. Carey School of Business
Department of Economics
Arizona State University
Tempe, AZ 85287-3806
Tempe, AZ 85287-3806
berthold.Herrendorf@asu.edu

Ákos Valentinyi
Economics, Arthur Lewis Building, 3rd floor
University of Manchester
Oxford Road
Manchester, M13 9PL
Manchester
United Kingdom
valentinyi.a@gmail.com

Richard Rogerson
Department of Economics
JRR Building Room 292
Princeton University
Princeton, NJ 08544
and NBER
rdr@princeton.edu

1 Introduction

Sustained growth in productivity is an essential part of development. The recent literature on structural transformation highlights that aggregate productivity depends on both *sectoral* productivity and *sectoral* composition.¹ During the early stages of development, sectoral composition importantly changes because a first phase of structural transformation takes place and labor moves from agriculture to other sectors. A popular notion is that successful development requires that the labor that leaves agriculture move into manufacturing rather than other non-agricultural sectors. In other words, successful development requires that the first phase of structural transformation involves industrialization.²

The literature suggests two channels through which expansion of manufacturing employment may be important for achieving aggregate productivity growth during this first phase of structural transformation. The first channel posits that manufacturing has a relatively high level of productivity, so moving labor to manufacturing has a positive level effect on aggregate productivity. The second channel posits that there is unconditional convergence in manufacturing productivity, so that productivity gaps relative to rich countries are narrowing in manufacturing and moving labor to manufacturing has a positive growth rate effect on aggregate productivity. If these two features uniquely describe manufacturing, then bypassing it and moving labor directly into services has detrimental effects on aggregate productivity.

Assessing the empirical relevance of the channels requires measures of sectoral productivity that are comparable across a large sample of poor countries. While data sets such as the Penn World Table (PWT) provide information about comparable productivity at the *aggregate* level of a large number of poor countries, similar information is not systematically available at the sectoral level.³ In this paper, we construct a data set of comparable labor productivity levels for agriculture and manufacturing for a large sample of mostly poor countries over the period 1990–2018.

Analysing our new dataset, we find only limited evidence in support of the two channels. Regarding the first channel, we find that while labor productivity gaps between rich and poor countries are smaller in manufacturing than in agriculture, labor productivity gaps in manufacturing are considerably larger than in the aggregate. That is, conditional on workers moving out of agriculture, aggregate gaps would decrease by more if the workers relocated to a sector other than manufacturing. Regarding the second channel, we find no evidence that labor productivity in manufacturing exhibits unconditional convergence. In fact, in our new dataset, the extent of

¹Herrendorf et al. (2014) provide a review of the overall literature on structural transformation. Gollin et al. (2002, 2007), Restuccia et al. (2008), Herrendorf and Valentinyi (2012), and Duarte and Restuccia (2020) are examples of analyses of sectoral productivity in the context of development.

²For examples of different versions of this view, see for example de Vries et al. (2015), Rodrik (2016), Diao et al. (2018), Diao et al. (2019), Lamba and Subramanian (2020), and Erumban and de Vries (2021).

³As we discuss in more detail later on, the Food and Agriculture Organization of the United Nations (FAO) has data on *agricultural* productivity levels in international prices, even though they stop in 1985.

convergence is somewhat lower in manufacturing than in the aggregate.

Turning to the details of our data construction, we leverage three recent data releases from the Groningen Growth and Development Center (GGDC): the Economic Transformation Database (ETD), the Africa Sector Database (ASD), and the Productivity Level Database (PLD). The ETD provides information on sectoral employment and sectoral value added, both nominal and real, for 51 countries during 1990–2018. We highlight two strengths of the ETD. First, it includes large representation of relatively poor countries from Africa, Asia, and Latin America. Second, the GGDC took extra care to harmonize the data on sectoral value added and employment, paying particular attention to the challenges arising in poor countries from the high incidence of small and informal firms and of the self-employed. We expand the coverage of the ETD by adding sectoral value added and employment data for 13 rich countries based on information from Eurostat, EUKLEMS, and other standard data sources. We label the resulting data set of 64 countries the Expanded Economic Transformation Database (EETD). It comprises covers all levels of development and all five continents, and it represents more than 4/5 of the world’s population and GDP in 2017.

Despite its many positive attributes, a limitation of the EETD is that it does not allow one to compare sectoral labor productivity levels across countries. That is, it does not contain the sectoral PPPs necessary to convert labor productivity levels measured in domestic prices to labor productivity levels measured in common international prices.⁴ Fortunately, the other two GGDC datasets, the ASD and the PLD, provide sectoral PPPs for the year 2005 that allow one to compute comparable measures of labor productivity at the sectoral level for a sample of 52 countries. While this sample includes representation from across the entire income distribution, it does not include 30 of the countries in the EETD, most of which are very poor.

We propose and validate a simple procedure that uses data from the combined ASD/PLD to impute agricultural and manufacturing labor productivity in 2005 in international prices for the countries in the EETD not covered by the ASD/PLD. Motivated by the notion that the law of one price approximately holds for tradables, our imputation procedure uses a log-linear regression of labor productivity levels in 2005 international prices on labor productivity levels in 2005 USD. We demonstrate that the regressions provide a very good fit for both agriculture and manufacturing throughout the entire income distribution of the ASD/PLD sample, with R-squared values that are close to one.

Having obtained measures of sectoral labor productivity in international prices for 2005 and the entire EETD sample, we compute sectoral labor productivity for all other years of the sample period 1990–2018 following standard practice. Specifically, in any year we calculate value added per worker in constant international prices from 2005 by using the domestic real growth rates of value added per worker reported by the EETD. Because our analysis focuses

⁴The PWT provides PPPs at the aggregate level, but they cannot be used to convert sectoral value added because relative prices vary systematically with the level of development [Sorensen (2001)].

entirely on labor productivity, in the interest of brevity we will treat productivity as synonymous with labor productivity in what follows.

Having constructed data on comparable sectoral productivity levels in agriculture and manufacturing for the 64 countries in the EETD, we use them to examine the two channels through which higher manufacturing employment may affect aggregate productivity in poor countries. Starting with the level effects, we investigate cross-country gaps in productivity levels at the sectoral and the aggregate level.⁵ We find that productivity gaps are considerably larger in agriculture than in the aggregate, which is to be expected given the evidence from the Food and Agricultural Organization (FAO) presented by Restuccia et al. (2008). While the productivity gaps in manufacturing are smaller than in agriculture, our first key result is that productivity gaps in manufacturing are actually *larger* than in the aggregate. These patterns imply that while there is some productivity gain in moving labor from agriculture to manufacturing, the gain could be larger if labor were to move to a sector other than manufacturing. Put differently, there is nothing special about moving workers from agriculture to manufacturing.

Turning to the growth-rate effects, our second key result concerns the convergence properties of sectoral and aggregate productivity. We consider both σ - and β -convergence.⁶ Time series plots for two different measures of cross-sectional dispersion – the 90-10 percentile gap in log productivity and the standard deviation of log productivity – reveal the same patterns: productivity dispersion has decreased modestly at the aggregate level, has remained fairly stable in agriculture, and has increased in manufacturing. That is, we do not find any evidence for σ -convergence in manufacturing. Similarly, a standard regression analysis does not produce evidence for unconditional β -convergence in either manufacturing or agriculture.

Our result of no convergence in manufacturing productivity over the period 1990–2018 mirrors the result in Inklaar and Timmer (2009) for an earlier period, but importantly their analysis was restricted to a sample of OECD countries. Our result sharply contrasts with the result in Rodrik (2013), who found rather strong degrees of unconditional convergence of manufacturing productivity over the period 1965–2005 using UNIDO data. We establish that an important reason for the starkly different results is that the GGDC measures of sectoral value added and employment are more comprehensive, addressing particularly the challenges arising from the high incidence of informal and small firms and of the self-employed in poor countries.

Our results largely challenge the view that expanding employment in manufacturing in the current set of poor economies is a recipe for successful development. This is of particular relevance for poor countries that are bypassing industrialization and instead undergoing what might be called service-led development; see Bah (2011) for examples. Our results suggest that

⁵We take comparable measures of aggregate productivity from the PWT.

⁶Loosely speaking, σ -convergence means that the standard deviation of the cross-country productivity distribution shrinks over time while β -convergence means that lower initial productivity levels are succeeded by higher productivity growth rates. See Barro and Sala-i-Martin (2003) for more details.

these countries are not necessarily doomed to experience disappointing aggregate productivity growth.⁷

Having found little evidence to support the notion that expanding manufacturing *employment* in today’s poorest economies is an important source of aggregate productivity growth, we end with investigating an alternative channel through which the manufacturing sector may be “special”. The alternative channel stresses manufacturing *productivity* and posits that achieving higher productivity growth in manufacturing is important because the same knowledge that allows for productivity growth in manufacturing at early stages of development will also promote productivity growth in other key sectors. While a definitive assessment of this channel is beyond the scope of this paper, we do present suggestive evidence in support of it. First, we show that whereas expansions in manufacturing employment have little correlation with aggregate productivity growth, manufacturing productivity growth is strongly correlated with aggregate productivity growth. Moreover, manufacturing productivity growth is also strongly correlated with productivity growth in several key non-agricultural sectors: trade, transportation, business services and finance. We conclude that it would be fruitful for future to study in more detail the origins of productivity growth in manufacturing and the other core sectors in poor countries.

The rest of the paper is organized as follows. In Section 2, we explain how we construct the new dataset. Section 3 examines productivity gaps and Section 4 examines productivity convergence. Section 5 analyzes the correlations among sectoral productivity growth. In Section 6, we conclude.

2 Sectoral Labor Productivity Measures

In this section, we introduce the three datasets from the GGDC that we leverage in our data construction. We then describe and verify our method for constructing cross-country measures of sectoral productivity levels for agriculture and manufacturing for a large sample of mostly poor countries. We also discuss to what extent our method applies outside of agriculture and manufacturing.

2.1 The Economic Transformation Database

The starting point for our analysis is the Economic Transformation Database of the GGDC. The ETD is a panel of 51 mostly poor countries in Asia, Africa and Latin America covering the period 1990–2018; see de Vries et al. (2021) for a detailed data description. It contains information on employment (persons engaged) and value added (both real and nominal in domestic

⁷Hallward-Driemeier and Nayyar (2017) discussed more generally how poor countries may achieve productivity growth outside of traditional manufacturing.

prices).⁸ The information is available for 12 sectors: agriculture, mining, manufacturing, construction, utilities, trade, transportation, financial services, business services, real estate, other services, and government. We note that in this classification the label of government should not be interpreted literally; in addition to the standard government activities of public administration, compulsory social security administration and defense, this category also includes education and human health and social work activities, independently of whether they are provided by the government.

An important strength of the ETD is that the GGDC harmonized the data collection across countries and devoted considerable effort to obtaining comprehensive measures of sectoral value added from national accounts and employment from labor force surveys and census data. The GGDC paid particular attention to the challenges arising from the high incidence of informal and own-account workers in poor countries. Despite its many positive features, an important limitation of the ETD is that it does not allow for the comparison of sectoral productivity levels across countries, because it does not provide sectoral PPPs with which one can calculate sectoral productivity levels in constant international prices. As we have discussed above, other GGDC data sets do provide this information for a subset of countries in the ETD. One contribution of our paper is to produce estimates of comparable sectoral productivity levels for the entire sample of countries in the ETD.⁹

Table 1 lists the 51 countries in the ETD. Upper-case letters indicate countries for which the GGDC reports sectoral PPPs elsewhere and lower-case letters indicate countries for which we will impute sectoral PPPs. We note each of Asia and Africa are represented by 21 countries, while Latin America is represented by 9 countries. As a result, the ETD is essentially a dataset of poor economies, containing only six rich economies: Hong Kong, Israel, Japan, Singapore, South Korea, and Taiwan.

Table 1: Countries in the Economic Transformation Database (ETD)

Africa	BOTSWANA	Burkina Faso	Cameroon	Egypt	
	ETHIOPIA	GHANA	KENYA	Lesotho	
	MALAWI	MAURITIUS	Morocco	Mozambique	
	Namibia	NIGERIA	Rwanda	SENEGAL	
	SOUTH AFRICA	TANZANIA	Tunisia	Uganda	ZAMBIA
Asia	Bangladesh	Cambodia	CHINA	Hong Kong	
	INDIA	INDONESIA	Israel	JAPAN	
	REP OF KOREA	Lao PDR	Malaysia	Myanmar	
	Nepal	Pakistan	Philippines	Singapore	
	Sri Lanka	Chinese Taipei	Thailand	TURKEY	Vietnam
	Latin America	ARGENTINA	Bolivia	BRAZIL	CHILE
Colombia		Costa Rica	Ecuador	MEXICO	

⁸The standard definition from the NIPA applies: persons engaged is the sum of full-time equivalent employees and the self-employed.

⁹As we will discuss in more detail below, data from the FAO have been used to generate sectoral productivity levels for agriculture, although the FAO series ends in 1985 and it is not clear how well it covers intermediate inputs. See Rao (1993) for further discussion.

To give a concrete idea of the economic characteristics of the countries in the ETD, Table 2 reports the distribution of GDP/worker for 1990 relative to the U.S., using GDP per worker data from the PWT.

Table 2: Distribution of GDP/Worker Relative to the U.S. (ETD, 1990)

Bins	< 0.05	[0.05, 0.10)	[0.10, 0.25)	[0.25, 0.50)	[0.50, 0.75)	[0.75, 1.00]	> 1.00
Countries	14	8	13	11	3	2	0

As the table shows, the sample of countries in the ETD is heavily skewed towards poor countries; more than forty percent of the countries have GDP/worker less than 10% of the value in the U.S. in 1990, and roughly two-thirds have GDP/worker less than 25% of the U.S. value. Only 5 countries, representing less than ten percent of the sample, have GDP/worker at least 50% of the U.S. value. The large representation of poor countries makes the ETD particularly attractive for studying productivity patterns in economies that are far away from the technological frontier.

The ETD also features a substantial degree of heterogeneity of growth experiences. Table 3 reports the distribution of the average annual growth rates of GDP/worker during the period 1990–2018 for all 51 countries of the ETD and for the 35 countries with initial GDP/worker less than 25% of the U.S. level.

Table 3: Distribution of Average Annual Growth Rates of GDP/Worker (ETD, in %)

Bins	< 1	[1, 2)	[2, 3)	[3, 4)	[4, 5)	> 5
All Countries	10	9	12	13	4	3
Poorest 35 Countries	5	7	6	10	4	3

The overall distribution of growth experiences is roughly uniform up to 4% per year, with a small but still significant right tail of countries with average growth rates exceeding 4% per year for almost three decades.¹⁰ The heterogeneity in growth experiences remains large even if we focus on the 35 poorest countries in our sample, with roughly half growing at more than 3% per year and half growing at less than 3% per year. The considerable heterogeneity in aggregate growth experiences will allow us to assess whether there are systematic patterns in sectoral productivity dynamics that vary with growth of GDP/worker.

¹⁰The seven high-growth-rate countries in the ETD are Brazil, China, Hong Kong, Israel, Laos, Myanmar, and Vietnam. While there are no countries from Africa in this group, we do note that there are 8 African countries with average annual growth rates greater than 3%.

2.2 Expanding the ETD

For some analysis it is useful to have greater representation from countries that are close to the frontier of aggregate productivity. We therefore expand the ETD sample by adding Australia, the U.S., and the following Western European countries: Austria, Belgium, Denmark, Finland, France, Germany, Italy, Netherlands, Spain, Sweden, and the U.K. The Australian data are from the OECD, the Swedish and the U.S. data are from the National Statistical Offices, and the data for the other European countries are from Eurostat and EUKLEMS. Fortunately, the GGDC reports sectoral PPPs for all of the additional rich countries.

We call the resulting database of 64 countries the Expanded Economic Transformation Database (EETD). For completeness, the country list of the EETD is in Appendix A.1. The EETD keeps the advantages of the ETD while considerably expanding coverage of countries to represent the entire world distribution of GDP per capita. To give concrete summary statistics, in 2017 the countries of the EETD represented more than 83% of the world's population and more than 82% of the world's GDP at international prices. Moreover, the EETD included 14 of the 15 most populous countries (with Russia missing) and 13 of the 15 largest economies (with Canada and Russia missing).

2.3 Existing Sectoral PPPs

Generating cross-country measures of sectoral productivity requires two inputs: measures of sectoral employment and measures of sectoral real value added based on a common set of prices.¹¹ We previously noted that one of the strengths of the ETD is the effort that went into measuring sectoral employment. The same is true for the sectoral data of the 13 additional rich countries that we added to the ETD to arrive at the EETD. In contrast, measuring levels of sectoral value added based on a common set of international prices remains a challenge.

Two relatively new datasets produced by the GGDC provide sectoral PPPs for 2005 for a sample of 52 countries, which allows for the calculation of sectoral value added in a common set of international prices of 2005. The Africa Sector Database (ASD) has sectoral PPPs for 11 Sub-Saharan African countries and the Productivity Level Database (PLD) has them for 42 countries.¹² The countries contained in the two datasets are listed in the Appendix A.1. Since South Africa features both in the ASD and the PLD, combining the two datasets gives sectoral PPPs for 52 countries in the year 2005. While the resulting 52-country sample does span a large range of GDP per capita levels, it has much less representation of poor countries and much more representation of rich countries than does the EETD.

¹¹Recall that we are using productivity to refer to labor productivity.

¹²de Vries et al. (2015) and Inklaar and Timmer (2014) provide more details on the two databases. Note that almost all sectoral PPPs reported by the GGDC are value-added PPPs. The exception is that the PLD reports gross-output PPPs instead of value-added PPPs for agriculture. We use whatever the GGDC offers.

We start with constructing comparable sectoral productivity levels in constant international prices of 2005 for the 52 countries of the combined ASD/PLD. To this end, we follow two standard steps. First, we calculate sectoral productivity in 2005 in current national prices by dividing sectoral value added in current national prices by sectoral employment. Second, we use the sectoral PPPs to convert the sectoral productivities from the first step into sectoral productivities for 2005 in international prices.

While the ASD/PLD collectively are an important step forward, their coverage of poor countries is still limited. Beyond the advanced economies of Japan and South Korea, the only Asian countries in the PLD are China, India and Indonesia, whereas the ETD has 15 additional Asian countries. Moreover, the PLD includes only 4 countries from Latin America – Argentina, Brazil, Chile and Mexico – whereas the ETD includes 9. Lastly, while the 11 Sub-Saharan African countries in the ASD provide some representation for Africa, the ETD has ten additional countries from Africa. In total, data on sectoral productivity levels is only available for 21 of the 51 countries in the ETD. The overlap between the ASD/PLD and the ETD was indicated in Table 1 with the use of upper-case letters.

2.4 Expanding the Coverage of Sectoral Productivity Levels

In this subsection, we produce comparable measures of productivity levels in agriculture and manufacturing for the countries in the EETD that are not covered by the ASD/PLD. The key challenge arises in step two of the previous calculation where we need a method to infer sectoral productivity in international prices for the subset of countries for which the ASD/PLD do not report sectoral PPPs. The issue of missing observations arises in many contexts, and we address it using standard methods. Specifically, we seek proxy variables that are both readily available for all countries and accurately predict sectoral productivities in international prices for the countries in the ASD/PLD. We estimate the relationship between sectoral productivities in international prices and the proxy variables on the sample of countries in the ASD/PLD and use it to impute sectoral productivities in international prices for the remaining countries in the EETD.

One tension in devising an imputation procedure is that expanding the set of proxy variables tends to increase in-sample predictive power but at the cost of over-fitting the data. Our implementation manages this tension by building on an insight of Rodrik (2013) that leads us to consider a single proxy variable and a linear predictor. He argued that because manufacturing output is tradable, the law of one price should approximately hold, implying that one can infer comparable cross-country productivity levels from manufacturing value added per worker expressed in United States dollars via the market exchange rate. While Rodrik focused on manufacturing, his insight should apply to other tradable sectors too. We therefore use it also for agriculture.

Rodrik (2013) used the identity function to predict manufacturing productivity in international prices from data on manufacturing productivity in USD. Importantly, he had neither the data to estimate a different relationship between the two productivity measures, nor could he validate that his assumed relationship was a reliable predictor. Our implementation addresses both issues by using the combined ASD/PLD to estimate the relationship between productivities in international prices and U.S. dollars (henceforth USD). We show that in 2005 a linear function of log productivities measured in USD provides an accurate prediction of log productivities measured in international prices. It is important to note that although the ASD/PLD has much less representation of poor countries than the ETD, the ASD/PLD still includes several poor countries. As a result, our projection is based on data throughout the entire cross-country income distribution.

While our data imputation leverages an insight from Rodrik (2013), we note several key differences between his procedure and ours. First, for any country of the EETD covered by the ASD/PLD, we use the measure of sectoral productivity generated by their sectoral PPPs in the benchmark year of 2005. Second for countries not covered by the ASD/PLD, rather than directly using productivity measured in USD in the benchmark year, we impute a productivity value for 2005 using the linear projection that we estimate on the data in the ASD/PLD.¹³ Third, concerning the calculation of productivity for years other than 2005, we follow standard procedure and generate the time series of comparable productivity levels by using the real sectoral productivity growth rates reported in the EETD. Our productivity levels are thus comparable both across countries and across time. In contrast, Rodrik (2013) used the manufacturing productivities in current USD as proxies for the manufacturing productivities in current international prices in all years, implying that his productivity levels were comparable only across countries in the same year but not over time.¹⁴

2.5 Constructing Productivity Levels for Tradables

In this subsection we show that log-linear functions of sectoral productivity in USD can be used to accurately estimate sectoral productivity in international prices for both agriculture and manufacturing in the benchmark year 2005. We denote 2005 productivities in country j measured in USD and in international prices by LP_j^{USD} and LP_j^{Int} , respectively, where it is understood that the productivities are at the level of either the agricultural or the manufacturing sector. Moreover, for each sector, we choose the units of international prices in 2005 such that

¹³As a practical matter this turns out to not be of particular significance; we find that the law of one price is a very good approximation for both manufacturing and agriculture in 2005, with the approximation being slightly better for manufacturing.

¹⁴One could make Rodrik's numbers comparable over time by applying the sectoral inflation in the U.S. to convert the value added in current USD prices into the value added in constant USD prices of 2005. Rodrik did not do this because it was not necessary for assessing the productivity convergence properties. Essentially, year fixed effects soak up a common inflation trend in the convergence regressions.

the values of productivity in USD and in international prices are equal for the U.S.:

$$LP_{US}^{USD} = LP_{US}^{Int}.$$

We then run the following regression for the 52 countries in the ASD/PLD:

$$\log LP_j^{Int} = \phi_0 + \phi_1 \log LP_j^{USD} + \varepsilon_j, \quad (1)$$

where ε_j is an error term. Rodrik's use of productivity measured in USD to directly proxy for productivity measured in international prices amounts to imposing that $\phi_0 = 0$ and $\phi_1 = 1$.

Table 4 shows the results of regression (1) for manufacturing. The value of ϕ_1 is quite precisely estimated and a 95 percent confidence interval includes 1. The point estimate of ϕ_0 is not statistically different from 0 at the five percent level. To interpret the economic significance of the point estimate of ϕ_0 one should compare it with the magnitude of the left-hand-side variable; since the mean value of log productivity is around 10, the point estimate of ϕ_0 effectively represents a second-order effect. The R-squared is 0.972.


Table 4: Manufacturing Productivities Regression (52 Countries in ASD/PLD, 2005)

ϕ_0	ϕ_1	R^2
0.377	0.960	0.972
(0.253)	(0.023)	

Robust standard errors are in parentheses.

Figure 1 presents the scatter plot of manufacturing productivity in USD and international prices, along with both the 45 degree line (solid) and the regression line (dashed). Both the regression results and the figure indicate a near perfect correlation between the 2005 manufacturing productivity in international prices and in USD. Importantly, the estimated relationship holds well throughout the income distribution. In particular, the three poorest countries in the sample – Ethiopia, Kenya and Malawi – all sit very close to the regression line. We conclude that for the ASD/PDL sample of 52 countries in 2005, data on manufacturing productivity in USD can be used to generate reliable estimates of manufacturing productivity measured in international prices via a simple linear predictor. We note that the results indicate that directly using USD productivity to proxy for productivity in international prices will also work well. Nonetheless, in what follows, we will use our regression estimates to construct productivity levels in 2005 for countries in the EETD not covered by the ASD/PLD.

Figure 1: Manufacturing Productivities in USD vs International Prices
(52 Countries in ASD/PLD, 2005)



Next, we repeat the analysis for the agricultural sector. Table 5 reports the results of running regression (1) for agriculture. The point estimate of ϕ_1 is again precisely estimated and a 95 percent confidence interval again includes 1. The point estimate of ϕ_0 is negative but is not statistically different from 0 even at the ten percent level. Given that the mean value of log productivity is almost ten, the magnitude of the point estimate makes it of second order economic importance. The R-squared value of 0.944 is again large, though a bit smaller than for manufacturing.

Table 5: Agricultural Productivities Regression (52 Countries in ASD/PLD, 2005)


ϕ_0	ϕ_1	R^2
-0.256 (0.327)	1.006 (0.035)	0.944

Robust standard errors are in parentheses.

Figure 2 shows a scatter plot of the data, along with both the 45 degree line (solid) and the regression line (dashed). We conclude that USD productivity can also be used to provide a very good approximation to agricultural productivity measured in international prices. While our regressions do not find statistical evidence against the law of one price, the fact that the R-squared is smaller for agriculture than for manufacturing does suggest that deviations from the law of one price are more prevalent in agriculture. This is consistent with a literature that does document sizeable trade and transportation barriers that increase domestic prices in agriculture in some countries; see for example Adamopoulos (2011) and Tombe (2015). Be that as it may, the extent of the deviations from the law of one price is small relative to the cross-sectional variation in productivity. Moreover, the deviations from the law of one price are not systematically larger in developing countries. Indeed, the correlation between log GDP per worker and differences in the two agricultural productivity measures, which capture the

deviations from the law of one price, is only 0.059 and is not significantly different from zero at the five-percent level.

Figure 2: Agricultural Productivities in USD vs. International Prices
(52 Countries in ASD/PLD, 2005)



In sum, the results of this subsection suggest that we can use productivity measured in USD to infer productivity in international prices for both manufacturing and agriculture in our benchmark year 2005. We will therefore apply this procedure to impute manufacturing and agricultural productivity levels for the 30 countries of the EETD for which the GGDC does not report sectoral PPPs. These countries tend to have low GDP per worker, so adding them gives a considerably greater coverage of poor countries than is provided by the combined ASD/PLD.

2.6 Productivity Levels of Goods, Market Services and Non-Market Services

The theoretical rationale for using productivity in USD to predict productivity in international prices for both manufacturing and agriculture derives from the idea that in the absence of trade barriers the law of one price applies to tradables. For completeness, it is of interest to examine the extent to which the law of one price does not hold for the non-tradable part of the economy. The combined ASD/PLD allows us to carry out an exercise of this sort for 52 countries, with the qualification that the available sectoral PPPs require us to consider a higher level of aggregation than the 12 sectors in the ETD. In what follows, we therefore split the economy into goods, market services, and non-market services. Goods comprise agriculture, construction, manufacturing, mining, and utilities; market services comprise trade, transportation, financial services, business services, and other services; and non-market services comprise real estate and government, which as noted earlier, includes education and healthcare. In terms of tradability, goods thus defined are still mostly tradable, though they now also contain the non-tradable components construction and utilities; market services includes some subsectors that are tradable,

while non-market services are the least tradable of the three sectors.


Table 6 shows the results of running the regression in (1) for the three-sector classification. Figure 3 presents the scatter plots, as well as the 45 degree line (solid) and the regression line (dashed).

Table 6: Goods and Services Productivities Regression (52 Countries in ASD/PLD, 2005)

	ϕ_0	ϕ_1	R^2
Goods	0.814 (0.292)	0.923 (0.027)	0.971
Market Services	4.237 (0.313)	0.605 (0.029)	0.922
Non-market Services	6.490 (0.657)	0.476 (0.061)	0.685

Robust standard errors are in parentheses.

Figure 3: Service Productivities in USD vs. International Prices
(52 Countries in ASD/PLD, 2005)



The results for goods are largely similar to what we found for agriculture and manufacturing, which is expected given that they dominate the goods sector. While the regression results provide statistical evidence against the law of one price for the overall goods sector, we note that the extent of the departure from the law of one price is minimal; the magnitude of the estimate for ϕ_0 is very small compared to the mean of log productivity, and while 1 lies outside of the 95 percent confidence bands for ϕ_1 , the upper band is still very close to 1.

In contrast, the results for the two services sectors differ quite markedly from our earlier results for agriculture and manufacturing. The point estimates of ϕ_1 are considerably smaller than one: 0.605 for market services and 0.476 for non-market services, and even a 99 percent confidence interval does not include 1. Additionally, the estimates of ϕ_0 are now statistically

different from 0 at the one-percent level, and the magnitude of the point estimates are now large compared to the mean levels of log productivity. Although the regression result for market services is evidence against the law of one price, the R-squared of 0.93 indicates that the productivity in USD could be used to impute productivity measured in international prices in the ASD/PLD sample. While this is perhaps of interest in a different context, we do not pursue it further in the current paper. In any case, the same does not apply to non-market services, where the R-squared is much smaller.

Our finding that the law of one price does not hold for services value added is closely related to the well-known Penn Effect, that is, in countries with low GDP per capita, the USD prices of final expenditures on services are lower than the international prices. Or put more casually, “services tend to be cheaper in poorer countries”. The Penn Effect was an important motivation behind the construction of the PWT because it implies that GDP per capita in poor countries increases when it is measured in international prices rather than in USD prices. We note that while our result is a relative of the Penn Effect, it is not the same. Our result is that a country with low productivity in a given services sector has a systematically low price in that services sector, whereas the Penn Effect references correlations with aggregate GDP per capita or worker. Moreover, and more importantly, our result pertains to the *value added* of the services sector whereas the Penn Effect pertains to the *final expenditure* on services. Since the final expenditure in a sector reflects value added from many other sectors, final expenditure and value added are distinct concepts with distinct properties.¹⁵

2.7 Comparison With FAO Data

As noted in the introduction, estimates of relative productivity of agriculture across countries were also generated using data from the FAO. The FAO provides measures of value added per worker in agriculture for 1970, 1975, 1980, and 1985 that are comparable across countries. These data have figured prominently in the macro development literature. For example, Restuccia et al. (2008) relied on them to argue that differences in agricultural productivity played a key role in shaping differences in aggregate productivity. Comparing the average of the top 5% of countries to the bottom 5% of countries, they find a ratio of 78 for agricultural productivity versus 5 for non-agricultural productivity. Combined with the fact that the poorest countries have most of their employment in agriculture, it becomes clear why agriculture is central to understanding the large differences in aggregate productivity across countries.

Given the prominence of the FAO data, it is of interest to compare our productivity measures for agriculture with those from the FAO. While the FAO data has attracted considerable attention, relatively little of this attention has been focused on the details of its construction, and it is important to review some of these details before proceeding with the comparison. Of particular

¹⁵Herrendorf et al. (2013) established how the distinction matters in the context of structural transformation.

significance is the fact that while the FAO produces measures of agricultural productivity for 103 countries, complete data are only available for 42 of them. Calculating productivity for the other 61 countries thus relies on imputation. The key step at which the imputation occurs is in moving from gross output to value added. In particular, while gross output and agricultural intermediate inputs are available for all countries, data on non-agricultural intermediate inputs (e.g., fertilizers, pesticides, energy) are only available for 42 countries. These 42 countries are used to estimate a mapping from gross output to value added, also utilizing data on other observables in their data set. The mapping is then used to generate value added data for the other 61 countries.

A key concern for assessing the reliability of an imputation procedure is whether the sample on which it is estimated adequately covers the range of countries for which it will be applied. We noted the same issue when describing our imputation procedure in Subsection 2.4. On this dimension there is cause for concern with the FAO imputation; only one of the 42 countries with a full complement of data is from Sub Saharan Africa while 24 are from European countries and their western offshoots. In contrast, the sample we used to estimate the imputation method included 11 countries from Sub Saharan Africa.¹⁶ The key message to take away from this discussion is that especially for lower income countries, the FAO measures should be understood as an imperfect proxy of productivity.

With the previous caveats in mind, we return to the comparison of our estimates with those resulting from the FAO data. We note two complicating features. The first is that the EETD starts in 1990 and the FAO data stop in 1985, which precludes comparing productivity levels for the same year. To the extent that productivity is relatively persistent, it is still of interest to compare levels across 1985 and 1990 for the subsample of 51 countries that are in both datasets; see Appendix A.1 for the list. The second complicating feature is that the international prices in the FAO data are from 1985 and in the GGDC are from 2005, implying that differences in productivity levels from the two datasets reflect differences in real productivity levels as well as differences in the international price levels. To minimize the effect of different benchmark years, we normalize U.S. agricultural labor productivity to one in both datasets and focus on comparing productivity levels relative to the U.S. across the two datasets.

Table 7 presents the results of regressing agricultural productivity relative to the U.S. in 1990 from the GGDC on agricultural productivity relative to the U.S. in 1985 from the FAO. The intercept is not significantly different from zero at the five percent level. While a 95 percent confidence interval for the slope does not include 1, it is not that far from 1.


¹⁶We also note that various measurement details differ between the FAO and the GGDC, including for example the treatment of inventories and farm housing expenditures.

Table 7: Agricultural Productivities Regression (51 Countries in $\text{FAO} \cap \text{EETD}$)

ϕ_0	ϕ_1	R^2
0.225	1.129	0.899
(0.142)	(0.044)	

Robust standard errors are in parentheses.

The scatter plot of the two productivity values in Figure 4 confirms the regression results: most points are close to the 45-degree line. We take this as confirmation that our measures track the earlier FAO numbers quite closely. This is of particular importance given that the FAO data stop in 1985 whereas our data provide coverage from 1990 through to 2018. Hence, our data allow us to track recent time series patterns in the cross-country distribution of productivity in agriculture that are no longer covered by the FAO data.

Figure 4: Agricultural Productivity in FAO and EETD (51 Countries in $\text{FAO} \cap \text{EETD}$)

While the two measures are highly correlated, we do want to draw attention to the fact that the countries in the bottom left corner of Figure 4 all lie below the 45 degree line, whereas the countries in the top right corner are relatively close to the 45 degree line. Importantly, this is not an artifact of our data imputation; this pattern holds even if we restrict attention to the subset of countries for which we do not do any imputation (i.e., the PLD/ASD sample). Given that many of the values for lower income countries in the FAO data involved imputation, we think it is reasonable to place more weight on the numbers derived from the GGDC data.

3 Cross-Country Productivity Gaps

Recent work on structural transformation and development has emphasized that the cross-country gaps in agricultural productivity are greater than the cross-country gaps in aggregate productivity, implying that moving labor from agriculture to non-agriculture in a poor economy

will necessarily close the aggregate productivity gap. This argument abstracts from the possibility that the effect may depend crucially on which sector within the non-agricultural part of the economy expands as a result of the contraction of the agricultural employment share, and in particular it does not say anything about the importance of moving the workers who leave agriculture into manufacturing versus some non-agricultural sector. Because our data construction yields comparable cross-country measures of productivity levels in both agriculture and manufacturing for a large sample of poor economies, we are uniquely positioned to shed light on the issue.

In what follows, we measure frontier productivity in a given sector and year by computing the average of the three highest productivity levels in that sector and year. Similarly, we compute frontier aggregate productivity by averaging across the three highest values of aggregate productivity. Gaps at the sectoral and aggregate level for each country are then defined as the log of their productivities relative to the frontier. Sectoral gaps are computed using our constructed data, while aggregate gaps are computed using data from the Penn World Table.

3.1 Productivity Gaps in Agriculture, Manufacturing, and the Aggregate

We start with the cross-country gaps in agricultural productivity versus those in aggregate productivity. Figure 5 shows the scatter plots in 1990, 2005 and 2018. Points below (above) the 45 degree line reflect countries for which the agricultural gap is larger (smaller) than the aggregate gap.

Figure 5: Agricultural vs. Aggregate Productivity Relative to Respective Frontier
(64 countries in EETD, 2005)





Figure 5 provides evidence for the common view that poor countries have considerably larger productivity gaps in agriculture than in the aggregate.¹⁷ The fact that we find larger

¹⁷The only outlier in this regard is Argentina, which has vast endowments with fertile land and so is closer to the frontier of agriculture than to the frontier of the aggregate.

productivity gaps in agriculture is consistent with previous work by Restuccia et al. (2008), who documented this pattern using the FAO data on agricultural productivity.¹⁸

That poor countries feature larger productivity gaps in agriculture than in the aggregate must mean that the productivity gaps in non-agriculture must be smaller than in the aggregate. However, to the best of our knowledge, there is no existing evidence that manufacturing has smaller productivity gaps than in agriculture or the aggregate. Figure 6 shows the scatter plots of manufacturing productivity gaps versus aggregate productivity gaps at the same three points in time – 1990, 2005 and 2018. Again, points below (above) the 45 degree line in these scatter plots reflect countries for which sectoral gaps are greater (smaller) than aggregate gaps.

Figure 6: Manufacturing vs. Aggregate Productivity Rel. to Respective Frontier
(64 countries in EETD)



Differently than for agriculture, the pattern for manufacturing shows significant changes over time. In 1990 the majority of points are below the 45 degree line, but are still relatively close to it. As time passes, the points tend to shift downward relative to the 45 degree line and in 2018 the pattern for manufacturing is quite similar to that in agriculture, with virtually all of the points for poor countries lying below the 45 degree line. Figure 6 suggests that while reallocating labor from agriculture to manufacturing would have reduced the aggregate productivity gap for a typical poor country in 1990, this is much less true in 2018.

The existing literature that analyzes agricultural productivity data from the FAO compares the size of productivity gaps in agriculture to the size of aggregate productivity gaps. For example, Caselli (2005) reports that the 90-10 ratio in agricultural productivity is 45 while the 90-10 ratio at the aggregate level is only 22.¹⁹ To better compare our estimates with these estimates, Table 8 reports the 90-10 percentile ratio in productivity levels in the aggregate, agriculture, and manufacturing for the 64 countries in the EETD in 1990, 2005, and 2018.

¹⁸Herrendorf and Valentinyi (2012) used the PWT to document that TFP gaps are larger in food production than in the aggregate, though their measure of food output was final expenditure rather than value added.

¹⁹Caselli also imputes productivity gaps for the entire non-agriculture sector and finds a 90-10 ratio of 4.

Table 8: 90-10 Percentile Productivity Gaps (64 countries in EETD)

Year	Aggregate	Agriculture	Manufact.
1990	24	94	26
2005	22	105	40
2018	17	101	42

Productivity gaps are ratios of productivity levels in constant international prices from 2005.

Three points stand out. First, consistent with the patterns that we saw in the figures, the gaps are largest in agriculture and smallest in the aggregate. Second, the gaps are modestly decreasing at the aggregate level, roughly constant in agriculture, and increasing in manufacturing. These patterns foreshadow the convergence results we will find in the next section. In particular, we will find very modest convergence in the aggregate and no evidence for convergence in either agriculture or manufacturing.

Third, and most striking, the magnitude of the sectoral gaps is very large: in 2018 the 90-10 ratio for manufacturing is more than twice the 90-10 ratio for the aggregate, and the ratio for agriculture is almost six times the ratio for the aggregate. As previously noted, our data imply even larger gaps than were found using FAO data, with the 90-10 ratio roughly twice as large as the number reported in Caselli (2005).²⁰ While our estimates suggest a significant upward revision to the size of agricultural productivity gaps, we view our estimates for manufacturing productivity gaps as more striking precisely because of the lack of previous evidence. Taking the 2018 gaps very literally, moving everyone in a poor economy into manufacturing would lead to aggregate gaps that are much larger than those currently observed. But the key point is that conditional on moving someone out of agriculture, our estimates imply that moving them into something other than manufacturing would have a much larger impact on aggregate productivity.

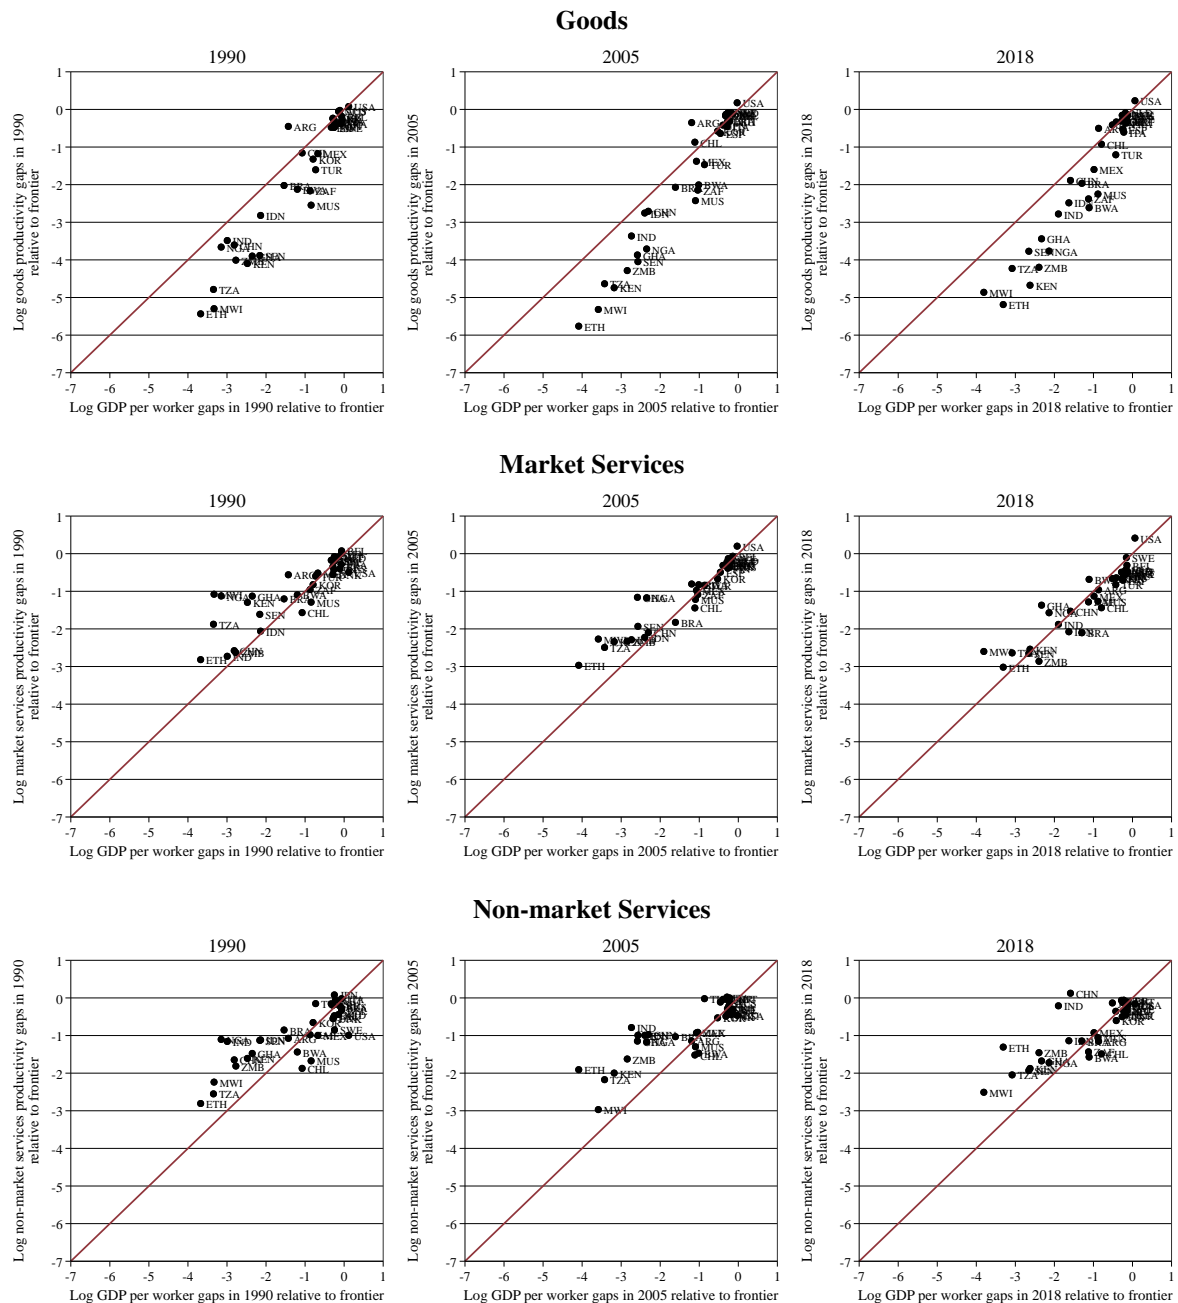
3.2 Productivity Gaps in Goods, Market Services and Non-Market Services

Since the sectoral gaps in agriculture and manufacturing are both larger than the aggregate gaps, it is also of interest to quantify the size of the gaps in the rest of the economy, most of which are services. To this end, we restrict attention to the subsample of the 34 countries that are in both the ASD/PLD and the EETD. Note that for this sample we do not need to do any imputation of sectoral productivities and have measures of relative sectoral productivity for a set of sectors with broader coverage. We use the same three-sector classification as earlier: goods, market

²⁰While a factor difference of 100 may seem implausibly large, we note that recent work by Gollin (2022) examines micro data from a small sample of rich and poor economies and finds that differences for his measure of output per hour can be larger than a factor of 1,000.

services and non-market services, and productivity gaps are again calculated relative to the appropriate frontier in each of 1990, 2005 and 2018. Figure 7 shows the results.²¹

Figure 7: Services vs. Aggregate Productivity Relative to Respective Frontier
(34 Countries in EETD \cap ASD/PLD)



As expected, the gaps in goods are larger than in the aggregate. More interesting, the gaps in *both* services sectors are lower than in the aggregate, with the *non-market* services

²¹We look at the common sample between EETD and ASD/PLD because we need the PPPs for 2005 to compute relative productivities in that year and the growth rates of real value added and employment to compute relative productivity in 1990 and 2018.

showing the smallest gaps. Pushed to the extreme, the mechanical reallocation logic would lead one to conclude that the largest improvement will arise if labor is reallocated to non-market services. One might discount this conclusion on the grounds that it is very challenging to measure productivity in the non-market services sector. But productivity gaps in market services are also smaller than aggregate gaps. So even if we disregard the non-market services sector, these observations offer a cautionary note about applying the mechanical reallocation logic to highlight the importance of industrialization to poor countries.

4 Productivity Convergence

The previous section found that while gaps in agricultural productivity have remained fairly stable, the gaps in manufacturing productivity have widened. This result is at odds with the influential work of Rodrik (2013), who found in UNIDO data that manufacturing productivity shows strong unconditional convergence during 1965–2005, and in particular for the subperiod 1995–2005. In this section, we examine the nature and extent of productivity convergence in the EETD in more detail, compare our results with those in Rodrik (2013), and discuss reasons for the differences.

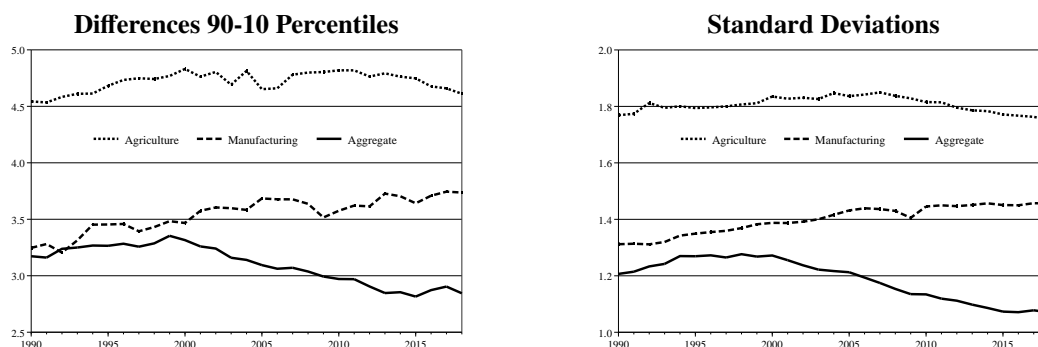
4.1 Convergence at the Sectoral and Aggregate Level

In this subsection, we examine σ - and β -convergence in agriculture, manufacturing and the aggregate for the 64 countries in the EETD. The novelty of our analysis is its focus on the convergence of *sectoral* productivity in manufacturing and agriculture; our *aggregate* convergence results add little to the literature and are included purely as a point of comparison with the sectoral convergence results for the EETD sample of countries and time period.²²

The concept of σ -convergence refers to a decrease in cross-sectional dispersion over time. To pursue σ -convergence, we consider two measures of dispersion: the 90-10 percentile difference in the log of productivity and the standard deviation of the log of productivity. Figure 8 shows the two measures for agriculture, manufacturing, and the aggregate over the period 1990–2018. The patterns foreshadowed by Table 8 are evident in both panels of the figure. First, the gaps are largest in agriculture, smallest in the aggregate, and intermediate in manufacturing. Second, and more importantly, while both dispersion measures for aggregate productivity decline significantly, there is little overall change for agriculture productivity and a considerably increase for manufacturing productivity. This contrasts starkly with the result of Rodrik (2013);

²²A large literature examines convergence of either GDP per capita or GDP per worker at the aggregate or sometimes regional level. (See, e.g., Barro (1991, 1997) and Barro and Sala-i-Martin (1992).) Two recent studies have revisited the convergence debate of the 1990s, extending the analysis to include the more recent decades: Acemoglu and Molina (2021) and Kremer et al. (1991).

Figure 8: Cross-country Dispersion in Log Productivity (64 countries in EETD, 2005)



he found that manufacturing exhibits σ -convergence, and in particular, to a greater extent than is found in the aggregate.

A second convergence concept refers to the tendency for countries with lower initial levels of productivity to have higher productivity growth rates. The standard approach for assessing such so called β -convergence is to regress the change in log productivity on the initial level of log productivity:²³

$$\Delta \log(LP_{jt}) = \alpha + \beta \log(LP_{jt-1}) + D_t + \varepsilon_{jt}, \quad (2)$$

where D_t are time fixed effects and LP_{jt} is productivity in constant international prices in country j and period t . The coefficient β measures the speed of *unconditional* convergence: if β is negative, then a lower initial productivity is associated with higher productivity growth and there is a tendency for low productivity countries to catch up to high productivity countries.²⁴

The literature also considers the specification in which country fixed effects D_j are added to the above regression. This allows the growth rate to depend on distance from a country specific level, and the coefficient β now measures the speed of *conditional* convergence: if β is negative, then each country converges to a productivity path that is determined by its institutional and economic characteristics captured by the country-fixed effects. Loosely speaking, one can view countries as converging to balanced growth paths that are parallel, with the level differences captured by the D_j .

Results for both specifications are presented in Table 9. The unconditional convergence results in columns (1), (3), and (5) are very similar for the aggregate, manufacturing, and agriculture: the point estimate of β is negative but very in absolute value and is not statistically different from zero at the five percent level. The point estimate for the aggregate is the largest in absolute value, but taking the point estimate at face value, the implied amount of convergence is minimal. If $\beta = -0.006$, a country with initial productivity that is 0.1 of the leader, will only

²³See Barro and Sala-i-Martin (2003) for a more detailed discussion of convergence regressions.

²⁴We note that unconditional β -convergence is necessary but not sufficient for σ -convergence.

Table 9: Convergence Regressions (64 countries in EETD, 1990–2018)

	Aggregate		Manufacturing		Agriculture	
	(1)	(2)	(3)	(4)	(5)	(6)
β	-0.006 (0.013)	-0.041 (0.011)	-0.003 (0.002)	-0.046 (0.014)	-0.002 (0.001)	-0.128 (0.023)
Number of observations			1,792			
Units	Constant international prices from 2005					
Time fixed effects	Yes					
Country fixed effects	No	Yes	No	Yes	No	Yes

Standard errors clustered at country level are in parentheses.

be at 0.143 of the leader after the 28 years of our sample period.²⁵ The absolute value of the point estimates for manufacturing and agriculture are less than half as large, implying that the extent of unconditional convergence is neither statistically nor economically significant.

The estimates for β in the conditional convergence specification in columns (2), (4), and (6) are again negative, but in contrast to the unconditional convergence specification, they are larger by an order of magnitude and are highly statistically significant. These estimates imply a substantial amount of conditional convergence over our sample period. For example, focusing on $\beta = -0.041$, which is the smallest point estimate in absolute value, and repeating the previous calculation, a country with productivity level of 0.1 relative to its balanced growth path level in 1990 will be at 0.490 of its balanced growth path level in 2018. The point estimate for agriculture, $\beta = -0.128$, implies even faster conditional convergence: a country that is at 0.1 relative to its balanced growth path productivity level in 1990 will be at 0.951 of its balanced growth level in 2018.

Just as growth experiences can vary across different samples of countries, the same may be true for convergence properties. One dimension of interest is geography; might countries in Africa, Asia and Latin America display different patterns of convergence? To pursue this we rerun our convergence specifications but leave out one of the following groups each time: Sub-Saharan Africa, South and East Asia, and Latin America. Note that we do not leave out the geographical area North Africa, Middle East and Central Asia because it has only four countries in the EETD – Egypt, Israel, Morocco, and Tunisia. Table 16 in the Appendix A.2 shows that this has very little effect on the estimates of β . We conclude that our β -convergence results for the overall sample are not hiding important heterogeneity across regions.²⁶

²⁵To obtain this number, subtract (2) for frontier productivity from (2) and take the exponential function:


$$\frac{LP_{t+28}}{LP_{t+28}^{\text{Front.}}} = \exp\left((1 + \beta)^{28} \log\left(\frac{LP_t}{LP_t^{\text{Front.}}}\right)\right) = 0.1^{(1+\beta)^{28}} \approx 0.143$$

²⁶We also run our convergence regressions for the subsample of 34 countries for which we not have to do any imputation of sectoral productivities in international prices. The results are very similar to those for the full sample; see Table 15 in Appendix A.2.

Having found some modest evidence for unconditional convergence in the aggregate, but effectively no evidence for unconditional convergence in agriculture and manufacturing, a natural follow-up question is whether the aggregate convergence is driven by convergence in some other sectors.²⁷ We cannot pursue this for our sample of 64 countries, but we can provide some answers for the sample of 34 countries which are both in the EETD and in the combined ASD/PLD. For these countries we have comparable sectoral productivity measures for a set of sectors with broader coverage. As in Section 3, we focus on the three-sector classification: goods, market services, and non-market services.

Figure 9 shows the graphs for σ -convergence. There is little evidence for σ -convergence in any of these three sectors, implying that the σ -convergence in the aggregate must have resulted from reallocation, most likely from the goods sector, which has higher dispersion, to the two services sectors, which have lower dispersion.²⁸

Figure 9: Cross-country Dispersion in Log Productivity of Goods and Services
(34 countries in $EETD \cap ASD/PLD$, 2005)



Turning to the results for β -convergence, Table 10 shows the regression results. The estimates for the unconditional convergence specification display greater statistical significance than the earlier results, and the point estimates for both services sectors are larger in absolute value than the earlier estimates for agriculture and manufacturing.²⁹ However, it remains the case that the point estimates imply a relatively minor extent of catch up over our 28 year period. In contrast, we find economically and statistically significant levels of conditional convergence for both goods and market services. The results for non-market services are mixed, with both a smaller point estimate in absolute value and less statistical significance.

In summary, the key finding from this subsection is that there is little to no evidence for

²⁷ Alternatively, the aggregate convergence could be driven by reallocation across sectors even in the absence of convergence at the sectoral level.

²⁸ Because dispersion differs between the two services sectors the reallocation could also have occurred between them. But as an empirical matter, with this three sector classification it is only the goods sector that experiences a declining employment share.

²⁹ As shown in the Appendix A.2, for the sample of 34 countries the point estimate for the regression using aggregate productivity is -0.005 and is significant at the five-percent level.

Table 10: Convergence Regressions for Goods and Services
(34 countries in EETD \cap ASD/PLD, 1990–2018)

	Goods		Market Services		Non-Mkt. Services	
	(1)	(2)	(3)	(4)	(5)	(6)
β	-0.002 (0.001)	-0.040 (0.014)	-0.009 (0.003)	-0.041 (0.013)	-0.007 (0.005)	-0.020 (0.022)
Observations				952		
Units			Constant international prices from 2005			
Time fixed effects			Yes			
Country fixed effects	No	Yes	No	Yes	No	Yes

Standard errors clustered at country level are in parentheses.

economically significant unconditional convergence in productivity of agriculture and manufacturing, or at the broader level of goods, market services and non-market services.

4.2 Comparing Results from GGDC and UNIDO Data

Our results for convergence in manufacturing productivity differ markedly from what Rodrik (2013) found: whereas we find practically no unconditional convergence in manufacturing productivity using data from the GGDC, he found strong unconditional convergence using data from UNIDO. In this subsection, we try to understand the reasons for the different results.

A first possible reason is differences in the construction of productivity series. As mentioned earlier, Rodrik computed productivity in every year as the value of productivity in current USD. In contrast, we use the USD measure to impute productivity in international prices only for 2005 and only for a subset of countries, and then use real productivity growth rates from domestic sources to fill in values for other years. In Appendix A.3 we show that this difference plays little role in accounting for the different results.

A second possible reason for the different results is that coverage differs starkly between the two datasets; whereas the GGDC seeks to provide a comprehensive measure of value added and employment that takes into account the importance of self-employment and informal enterprises in poor countries, the UNIDO data are based on surveys that restrict coverage to relatively large firms that operate in the formal sector.³⁰ Rodrik (2013) openly acknowledged the possibility that the limitations of the UNIDO data might affect the scope of his finding. In particular, he wrote in his abstract: *“Because of data coverage, these findings should be as viewed as applying to the organized, formal parts of manufacturing.”* Unfortunately, many papers that cite the findings in Rodrik (2013) do not seem to acknowledge this important qualification.

³⁰In fact, there are additional differences beyond data coverage. The GGDC measurement of value added adopts methods used in constructing NIPA accounts, whereas the UNIDO data are based on surveys that do not necessarily adhere to the same accounting practices. We are unable to say anything about the effect of these differences.

In the rest of this section we show that the difference in coverage is an important reason for the differences in results, and that in many poor countries the scope of Rodrik’s result is rather narrow.

4.2.1 The Role of Different Underlying Data Sources

Our convergence results using GGDC data are for a different time period and a different set of countries than were Rodrik’s convergence results using UNIDO data. Because convergence properties can be specific to the set of countries and the time period, it is important to compare results from the two datasets for a common time period and sample of countries. With this in mind we focus on the period 1995–2005 and on the 41 countries that are present in both the EETD and the UNIDO data sets.³¹ For both the GGDC and the UNIDO data, we follow the procedure used by Rodrik (2013) and construct productivity measured in current USD in every year. We then run the convergence regression (2) on these two datasets. Because we are interested in unconditional convergence Table 11 reports results only for the specification without country fixed effects.

Table 11: Convergence Regressions for Manufacturing in Current USD Prices, EETD versus UNIDO (41 countries in EETD \cap UNIDO, 1995–2005)

	EETD	UNIDO
β	–0.007 (0.005)	–0.020 (0.006)
Number of observations	410	
Units	Current prices in USD	
Time fixed effects	Yes	
Country fixed effects	No	

Standard errors clustered at country level are in parentheses.


Strikingly, the conclusions about unconditional convergence differ quite dramatically between the GGDC and the UNIDO data. Whereas the point estimate of β is both small (–0.007) and not statistically different from zero at the five-percent level when using GGDC data, the estimate is –0.020 and significant at the one percent level when using the UNIDO data. That is, holding fixed the time period and set of countries considered, using productivity data from the GGDC instead of from UNIDO dramatically changes the conclusions regarding unconditional convergence.³²

³¹The UNIDO data we use in this section are INDSTAT2, ISIC Revision 3. A list of the overlapping countries is contained in Appendix A.1.

³²Repeating our earlier calculation, $\beta = -0.02$ implies that a country that begins at 0.1 of the leader will be at 0.27 of the leader after 28 years.

The different properties of the GGDC and UNIDO datasets are also apparent if we plot time series changes in two measures of cross-sectional dispersion used before, namely, the 90-10 ratio and the standard deviation of log manufacturing productivity over the period 1995–2005. Figure 10 shows a marked decline for UNIDO data and a marked rise for the GGDC data.

Figure 10: Cross-country Dispersion in Manufacturing Productivity
(41 countries in $EETD \cap UNIDO$, 1995–2005)



4.2.2 Documenting Differences in Coverage

We have already highlighted that the EETD and UNIDO datasets provide different coverage, and there seems to be no controversy about this. In this subsection, we interpret the EETD data as a measure for the overall manufacturing sector and the UNIDO data as a measure for a particular segment of the manufacturing sector, and we document the size and nature of the difference in coverage and note some implications. To best focus on the set of poorer countries that are of greatest interest, we restrict attention to the overlapping samples between UNIDO and the ETD for the period 1990–2018. This results in a sample of 30 countries.³³

For each year and country we calculate the ratio of manufacturing employment in the UNIDO data to manufacturing employment in the GGDC data. We call this ratio the *coverage ratio*. We then average the values across all years for each country. Table 12 shows the distribution of values.

Table 12: Coverage Ratios UNIDO–EETD Manufacturing Employment
(30 countries in $EETD \cap UNIDO$, 1990–2018)

UNIDO Employment / EETD Employment	0–0.25	0.25–0.50	0.50–0.75	0.75–1.00
Number of Countries	2	11	11	6

For more than 40% of the countries that are contained in both datasets, the UNIDO data cover less than half of the employment in manufacturing that is reported by the GGDC data.

³³There are a few countries in the UNIDO dataset that have missing observations for some years in the period 1990–2018. We include countries as long as they contain data for the majority of years and have observations both in the beginning and end of the sample. We fill in missing observations using linear interpolation.

For these countries the patterns in the UNIDO data would only be indicative of overall patterns if the behavior of the uncovered segment was very similar to that of the covered segment. However, there is good reason to view this as being unlikely; even in advanced countries we know that productivity dynamics differ quite markedly across the firm size distribution. We note that the smaller the size of the covered segment, the more severe this selection issue becomes.

The large cross-sectional variation in the extent of coverage implies that any cross-country comparison reflects a mixture of “true” differences as well as differential selection effects. This makes it very challenging to interpret the size of cross-sectional differences in a given year. An additional issue arises because some countries experience significant changes in the coverage ratio over time. Figure 11 shows a scatter plot of values in 1990 versus 2018. While the points do tend to follow the 45 degree line (the correlation is 0.52), the figure also indicates that there are large deviations, both up and down, for individual countries.

Figure 11: Changes in the Manufacturing Employment Coverage Ratios UNIDO–EETD (30 countries in $EETD \cap UNIDO$)





Figure 12 illustrates the changes for four populous countries: Kenya, Mexico, the Philippines and Tanzania. Kenya and Tanzania experience large drops, and as of 2018 the UNIDO measure is missing almost 90% of total employment in Kenyan manufacturing. In contrast, Mexico displays a large increase in coverage. The increases in coverage in Mexico are largely captured by two discrete jumps, and although not shown, these jumps reflect changes in the UNIDO measure, which most probably reflects discrete changes in the coverage of the UNIDO survey. Changes in coverage ratios over time make it challenging to interpret time series changes in cross-sectional measures of dispersion.

In fact, Figure 13 shows that there is a systematic negative relationship between the productivity growth measure in UNIDO and the change in our coverage ratio.³⁴ The figure reinforces our earlier point about interpreting cross-country differences in productivity levels. Because


³⁴Nepal is a strong outlier. Leaving it out, the slope is statistically significantly different from zero.

Figure 12: Manufacturing Employment Ratio in UNIDO and EETD (Selected Countries)



selection effects are potentially changing over time, it is extremely difficult to interpret convergence results using UNIDO data.

Figure 13: Manufacturing Productivity Growth in UNIDO versus Change in Coverage Employment Ratio (30 countries in EETD \cap UNIDO, 1990–2018)



4.2.3 Summary and Discussion

The preceding analysis points to differences in coverage between the UNIDO and the GGDC data as playing an important role in accounting for the different patterns of convergence properties over the 1995–2005 period. If one is interested in characterizing the properties of the overall manufacturing sector, the results using GGDC data seem to be the relevant ones. However, to better understand the source of aggregate differences, it may also be of interest to systematically compare specific segments of the manufacturing sector across countries. For such a comparison to make sense, the segments must be defined systematically and in a consistent way across

countries. Unfortunately, the UNIDO data do not employ a harmonized definition for large or formal part of manufacturing, neither across countries at a point in time nor within a given country over time.

While the work of Rodrik (2013) represents a first step in assessing cross-country differences in productivity dynamics within particular segments of the manufacturing sector, and it uncovered an interesting pattern, the limitations of UNIDO data make it difficult to interpret his results. Additional work is needed to draw more definitive conclusions. We think the work with micro data in the recent paper by Diao et al. (2001) is an example of the type of research that is needed to make progress in documenting the differential behavior of formal and informal firms.

5 Additional Evidence and The Significance of Manufacturing Productivity Growth

So far, we have focused on two possible channels through which reallocating workers into manufacturing in poor economies could close the aggregate productivity gap between them and advanced economies: (i) gaps in manufacturing productivity could be smaller than gaps in the aggregate productivity and (ii) productivity growth in manufacturing could be catching up to the frontier faster than in other sectors. Our analysis of the sectoral productivity data has suggested that the two channels are of limited empirical relevance for today's poor economies.

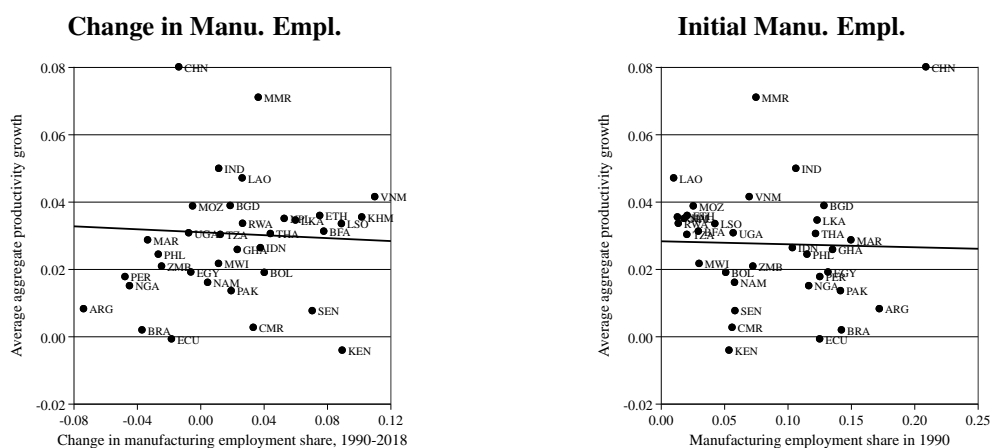
In the first part of this section we present additional evidence against the view that manufacturing employment has a key role to play in allowing poor countries to close the aggregate productivity gap with advanced countries. In particular, we leverage the heterogeneity in growth experiences among the poor countries in the EETD to show that higher manufacturing employment is not positively related to aggregate productivity growth. In the second part of this section we present evidence that is consistent with the manufacturing sector playing an important role in the development of today's poor economies even if manufacturing employment does not. In particular, we present evidence showing that the growth rate of manufacturing productivity is highly positively correlated with aggregate productivity growth.

To focus the attention on the countries in the earliest stages of development, our analysis in this section uses the EETD subsample of the 35 countries with initial GDP/worker less than 25% of the U.S. level. We have previously studied the same subsample in Subsection 5.1. As noted there, the heterogeneity in the aggregate productivity growth experiences within this sample provides the opportunity to uncover systematic patterns that connect sectoral and aggregate productivity growth.

5.1 Manufacturing Employment and Aggregate Productivity Growth

If moving workers into manufacturing was an important driver of overall productivity growth in the early stages of development, then one would expect to see a high positive correlation between changes in the manufacturing employment share and aggregate productivity growth for our sample of 35 relatively poor countries. But the value of this correlation is -0.25 and is not significantly different from zero at the five-percent level. And if we leave out China, which is an outlier, the correlation increases, but is virtually zero at -0.02 . The left panel of Figure 14 shows a scatter plot of the two variables. Alternatively, if the absolute size of the manufacturing sector is what matters, it may be more informative to consider the correlation between the initial employment share in manufacturing and aggregate productivity growth. But the value of this correlation is -0.03 , and again it is not significantly different from zero at the five-percent level. The right panel of Figure 14 shows the scatter plot of the two variables.

Figure 14: Manufacturing Employment versus Aggregate Productivity Growth
(35 poor countries in EETD, 1990–2018)




While we emphasize that these are just raw correlations, they provide an additional, simple challenge to the notion that the manufacturing employment share plays an important role in generating overall productivity growth.³⁵

5.2 Productivity Correlations

We begin by examining the relationship between manufacturing and aggregate productivity growth for our sample of 35 poor countries. Figure 15 presents a scatter plot of the data.

³⁵We obtain very similar correlations if we break 1990–2018 into two subperiods and consider each separately. If we regress average aggregate productivity growth on both the initial level and the change in the manufacturing employment share then neither variable is significant.

Figure 15: Manufacturing versus Aggregate Productivity Growth
(35 poor countries in EETD, average of 1990–2018)



The scatter plot reveals a strong positive relationship and the correlation is 0.65. This is in striking contrast to the very low correlations we found between aggregate productivity growth and the two measures of manufacturing employment in the previous subsection, and it suggests that it could be useful to consider the role of manufacturing productivity *growth* in the early development process. Is the high positive correlation between manufacturing and aggregate productivity growth a distinguishing feature of the manufacturing sector relative to other sectors? To answer this question, Table 13 reports the same correlation for each sector.

Table 13: Correlation Between Sectoral and Aggregate Productivity Growth

Agri.	Mining	Manu.	Constr.	Util.	Transp.	Trade	Bus.	Finan.	Oth. Serv.	Gov't
0.33	0.50	0.65	0.37	0.41	0.73	0.62	0.52	0.54	0.51	0.64

While manufacturing stands out as having one of the highest correlations, there are others with similar values, and transportation has an even higher correlation. To proceed further, we examine the correlation patterns for productivity growth across sectors. Table 14 focuses on

Table 14: Productivity Correlations

	Agricult.	Manufact.	Trade	Gov't
Agricult.	1.00	0.27	0.08	0.05
Mining	-0.05	0.55	0.40	0.43
Manufact.	0.27	1.00	0.69	0.43
Construct.	-0.02	0.61	0.44	0.14
Utilities	0.14	0.54	0.45	0.12
Transp.	0.10	0.66	0.66	0.48
Trade	0.08	0.69	1.00	0.37
Business	0.16	0.69	0.57	0.43
Finance	0.16	0.68	0.46	0.38
Other Serv.	0.14	0.22	0.40	0.23
Gov't	0.05	0.43	0.37	1.00

these correlations for each of the four largest sectors: agriculture, manufacturing, trade and government. The table reveals interesting differences across the four largest sectors. First, the correlations between agriculture and the other sectors are very low indeed: the highest correlation is with manufacturing and is only 0.27; all of the other correlations are less than 0.20. Second, government productivity growth displays higher correlations with the other sectors than does agriculture, but all of the correlations are below 0.50 and several are less than 0.25. The relative independence of productivity growth in both agriculture and government of productivity growth in the other sectors suggests that the additional know-how that generates productivity growth in the two sectors has little benefit in terms of lifting productivity elsewhere in the economy. Conversely, it also suggests that the additional know-how that generates productivity growth elsewhere in the economy is of relatively little value in these two sectors. Note that despite the lack of any “spillover” benefits, productivity growth in agriculture and government may still have important aggregate effects given that sectoral employment shares tend to be large in poor countries.

Continuing with productivity growth in manufacturing and trade in Table 14, the situation is very different. In particular, productivity growth in trade features higher correlations than government with productivity growth in every sector, with the lone exception of mining. Three of the correlations are above 0.5 and two are above 0.6. This pattern intensifies for manufacturing; manufacturing productivity growth exhibits strong positive correlations with productivity growth in several other sectors, with five values above 0.50 and four values above 0.60. In fact, manufacturing has the highest average correlation of all sectors, not just of the four large sectors shown here. Moreover, as evidence of the somewhat special nature of manufacturing, a large majority of sectors display their highest correlation with the manufacturing sector.

More generally, if one looks at the full set of bivariate correlations across sectors, there are only four sectors which feature in connections with correlations that exceed 0.60. These four sectors are manufacturing, trade, transportation and business services. Looking at the full set of bivariate correlations among these four sectors, all of the correlations save one are greater than 0.6. Importantly, as Table 13 shows, three of these four sectors feature the highest correlations between their own productivity growth and aggregate productivity growth: manufacturing, transportation and trade with correlations of 0.65, 0.73 and 0.62, respectively.

5.3 Discussion

The preceding analysis points to a group of four sectors that stand out in terms of their connection to aggregate productivity growth: manufacturing, trade, transport and business services. This result is consistent with a simple and intuitive narrative. The early stage of development are largely about moving labor out of subsistence agriculture and into other activities. Achieving sustained improvements in aggregate productivity during the early stage of development

requires productivity growth in the activities that are absorbing labor. This requires a transition from a state with relatively little exchange of goods and services to one in which specialization and exchange of goods and services increase dramatically. The core absorbing activities in this context will be those that reflect the production and distribution of goods: manufacturing, trade, transportation and business services. Therefore, achieving high levels of aggregate productivity growth very much depends on achieving high levels of productivity growth in these four sectors.

The high correlation of productivity growth among sectors could be caused by spillovers. Focusing only on the four sectors, it is manufacturing that displays the highest average correlation with the others, and in each case the highest bivariate correlation for each of trade, transportation and business services is with manufacturing. While these correlations are consistent with productivity growth in manufacturing causing productivity growth in other sectors, they are also consistent with one or more of the other sectors being the driving force.

Alternatively, the high correlation of productivity growth among these sectors could indicate a common driving force. In particular, we think that it is plausible that the additional know-how needed to generate productivity gains in these four sectors in the early stages of development is not particularly specific to any individual sector. Viewed from this perspective, the evolution of manufacturing productivity growth may simply serve as a barometer for the extent to which the economy is accumulating the general know-how that enables productivity growth in these four sectors. Whereas an economy could logically experience high productivity growth by increasing productivity in all activities except manufacturing, the data indicate that this is simply not a typical path.

A definitive interpretation of these productivity correlations is beyond the scope of our paper and will require much more data than is available in the GGDC databases in general and the EETD in particular. The correlation evidence we have reported establishes neither causation nor its direction. Nonetheless, we think that the correlation evidence is an interesting first step, which, at the very least, is suggestive of a special role for manufacturing productivity.

6 Conclusion

We have leveraged recent data from the GGDC to build a new dataset of comparable productivity levels in agriculture and manufacturing for 64 mostly poor countries during 1990–2018. Our analysis has been purely empirical, but the data that we have constructed will be an important input into model based analyses that seek to understand the forces that shape structural change during the early phases of development and the heterogeneous experiences across early developers.

We have used the new dataset to assess the popular notion that industrialization – the move-

ment of labor into manufacturing – plays an important role in helping poor countries close productivity gaps with advanced economies. We have found only limited support for the popular notion, suggesting that service-led development may be much less of an issue than is often thought. An example where this conclusion applies is India, which authors like Kochhar et al. (2006) or Lamba and Subramanian (2020) have bemoaned for its lack of industrialization.

We have also used the new dataset to investigate the possibility that manufacturing sector may be special because the same knowledge that allows for productivity growth in manufacturing at early stages of development also promotes productivity growth in other key sectors. While a definitive assessment of this channel is beyond the scope of this paper, we have presented suggestive evidence in support of it. An important avenue for future research is to investigate what the determinants of the productivity growth patterns in the core sectors are.

References

- Acemoglu, Daron and Carlos A. Molina**, “Converging to Converge? A Comment,” Working Paper 28992, NBER, Cambridge, MA 2021.
- Adamopoulos, Tasso**, “Transportation Costs, Agricultural Productivity, and Cross-Country Income Differences,” *International Economic Review*, 2011, 52, 489–521.
- Bah, El-hadj**, “Structural Transformation Paths Across Countries,” *Emerging Markets Finance and Trade*, 2011, 47, 5–19.
- Barro, Robert J. and Xavier Sala-i-Martin**, “Convergence,” *Journal of Political Economy*, 1992, 100, 223–251.
- and —, *Economic Growth*, 2nd ed., New York: MIT Press, 2003.
- Caselli, Francesco**, “Accounting for Cross-Country Income Differences,” in Philippe Aghion and Steven Durlauf, eds., *Handbook of Economic Growth*, Vol. 1A, Amsterdam and New York: North Holland, 2005, chapter 9, pp. 679–742.
- de Vries, Gaaitzen, Linda Arfelt, Dorothea Drees, Mareike Godemann, Calumn Hamilton, Bente Jessen-Thiesen, Ahmet Ihsan Kaya, Hagen Kruse, Emmanuel Mensah, and Pieter Woltjer**, “The Economic Transformation Database (ETD): Content, Sources, and Methods,” Technical Note, WIDER 2021.
- de Vries, Gaaitzen, Marcel Timmer, and Klaas de Vries**, “Structural Transformation in Africa: Static Gains, Dynamic Losses,” *Journal of Development Studies*, 2015, 51, 674–688.

- Diao, Xinshen, Margaret McMillan, and Dani Rodrik**, “The Recent Growth Boom in Developing Economies: A Structural-Change Perspective,” in Machiko Nissanke and José Antonio Ocampo, eds., *The Palgrave Handbook of Development Economics*, Cham: Palgrave Macmillan, 2019, chapter 9, pp. 281–334.
- , —, and **Sam Wangwe**, “Agricultural Labour Productivity and Industrialisation: Lessons for Africa,” *Journal of African Economics*, 2018, 28, 28–65.
- , **Mia Ellis, Margaret S. McMillan, and Dani Rodrik**, “Africa’s Manufacturing Puzzle: Evidence from Tanzanian and Ethiopian Firms,” Working Paper 28344, NBER 2001.
- Duarte, Margarida and Diego Restuccia**, “Relative Prices and Sectoral Productivity,” *Journal of the European Economic Association*, 2020, 18, 1400–1443.
- Erumban, Abdul and Gaaitzen de Vries**, “Industrialization in Developing Countries: Is It Related to Poverty Reduction?,” Working Paper 172, WIDER 2021.
- Gollin, Douglas**, “Farm Size and Productivity: Lessons from Recent Literature,” mimeo, Oxford University 2022.
- , **Stephen L. Parente, and Richard Rogerson**, “The Role of Agriculture in Development,” *American Economic Review, Papers and Proceedings*, 2002, 92, 160–164.
- , —, and —, “The Food Problem and the Evolution of International Income Levels,” *Journal of Monetary Economics*, 2007, 54, 1230–1255.
- Hallward-Driemeier, Mary and Gaurav Nayyar**, *Trouble in the Making? The Future of Manufacturing-Led Development*, Washington DC: World Bank Publications, 2017.
- Herrendorf, Berthold and Ákos Valentinyi**, “Which Sectors Make Poor Countries so Unproductive?,” *Journal of the European Economic Association*, 2012, 10, 323–341.
- , **Richard Rogerson, and Ákos Valentinyi**, “Two Perspectives on Preferences and Structural Transformation,” *American Economic Review*, 2013, 103, 2752–2789.
- , —, and —, “Growth and Structural Transformation,” in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth*, Vol. 2, Elsevier, 2014, pp. 855–941.
- Inklaar, Robert and Marcel P. Timmer**, “Productivity Convergence across Industries and Countries: The Importance of Theory-Based Measurement,” *Macroeconomic Dynamics*, 2009, 13, 2188–240.

— **and** —, “The Relative Price of Services,” *Review of Income and Wealth*, 2014, 60, 727–746.

Kochhar, Kalpana, Utsav Kumar, Raghuram Rajan, Arvind Subramanian, and Ioannis Tokatlidis, “India’s Pattern of Development: What Happened, What Follows?,” *Journal of Monetary Economics*, 2006, 53, 981–1019.

Kremer, Michael, Jack Willis, and Yang You, “Converging to Convergence,” in Martin S. Eichenbaum and Erik Hurst, eds., *NBER Macroeconomic Annual*, Vol. 36, Cambridge, MA: MIT Press, 1991.

Lamba, Rohit and Arvind Subramanian, “Dynamism with Incommensurate Development: The Distinctive Indian Model,” *Journal of Economic Perspectives*, 2020, 34, 3–30.

Rao, D.S. Prasada, “Intercountry Comparisons of Agricultural Output and Productivity,” FAO Economic and Social Development Series 112, Food and Agriculture Organization of the United Nations 1993.

Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu, “Agriculture and Aggregate Productivity: A Quantitative Cross–Country Analysis,” *Journal of Monetary Economics*, 2008, 55, 234–50.

Rodrik, Dani, “Unconditional Convergence in Manufacturing,” *Quarterly Journal of Economics*, 2013, 128, 165–204.

—, “Premature Deindustrialisation,” *Journal of Economic Growth*, 2016, 21, 1–33.

Sorensen, Anders, “Comparing Apples to Oranges: Productivity Convergence and Measurement across Industries and Countries: Comment,” *American Economic Review*, 2001, 91, 1160–1167.

Tombe, Trevor, “The Missing Food Problem: Trade, Agriculture, and International Productivity Differences,” *American Economic Journal: Macroeconomics*, 2015, 7, 226–258.

A Appendix

A.1 Country Lists

Countries in the EETD

Argentina; Australia; Austria; Bangladesh; Belgium; Bolivia; Botswana; Brazil; Burkina Faso; Cambodia; Cameroon; Chile; China; Colombia; Costa Rica; Denmark; Ecuador; Egypt; Ethiopia;

Finland; France; Germany; Ghana; Hong Kong; India; Indonesia; Israel; Italy; Japan; Kenya; Rep. of Korea; Lao PDR; Lesotho; Malawi; Malaysia; Mauritius; Mexico; Morocco; Mozambique; Myanmar; Namibia; Nepal; Netherlands; Nigeria; Pakistan; Peru; Philippines; Rwanda; Senegal; Singapore; South Africa; Spain; Sri Lanka; Sweden; Tanzania; Chinese Taipei; Thailand; Tunisia; Turkey; Uganda; United Kingdom; United States; Vietnam; Zambia

ASD

Botswana; Ethiopia; Ghana; Kenya; Malawi; Mauritius; Nigeria; Senegal; South Africa; Tanzania; Zambia.

PLD

Argentina; Australia; Austria; Belgium; Brazil; Bulgaria; Canada; Chile; China; Cyprus; Czech Republic; Denmark; Estonia; Finland; France; Germany; Greece; Hungary; India; Indonesia; Ireland; Italy; Japan; Latvia; Lithuania; Luxembourg; Malta; Mexico; Netherlands; Poland; Portugal; Romania; Russia; Slovakia; Slovenia; South Africa; South Korea; Spain; Sweden; Turkey; United Kingdom; United States.

Overlap of EETD and FAO Databases

Argentina; Australia; Austria; Bangladesh; Belgium; Bolivia; Brazil; Burkina Faso; Cameroon; Chile; Colombia; Costa Rica; Denmark; Ecuador; Egypt; Ethiopia; Finland; France; Germany; Ghana; India; Indonesia; Israel; Italy; Japan; Kenya; Malawi; Malaysia; Mexico; Morocco; Mozambique; Nepal; Netherlands; Nigeria; Pakistan; Peru; Philippines; Republic of Korea; Rwanda; Senegal; South Africa; Spain; Sri Lanka; Sweden; Tanzania; Thailand; Tunisia; Turkey; Uganda; United Kingdom; United States.

Overlap of EETD and UNIDO Databases 1995–2005

Australia; Austria; Belgium; Bolivia; Brazil; China; Colombia; Costa Rica; Denmark; Ecuador; Egypt; Ethiopia; Finland; France; Ghana; Hong Kong; India; Indonesia; Israel; Italy; Japan; Kenya; Malawi; Malaysia; Mauritius; Mexico; Morocco; Netherlands; Peru; Philippines; Republic of Korea; Senegal; Singapore; South Africa; Spain; Sweden; Chinese Taipei; Thailand; Turkey; United Kingdom; United States.

EETD Countries with GDP pc less than 25% of the U.S. in 1990

Argentina; Bangladesh; Bolivia; Brazil; Burkina Faso; Cambodia; Cameroon; China; Ecuador; Egypt; Ethiopia; Ghana; India; Indonesia; Kenya; Laos; Lesotho; Malawi; Morocco; Mozam-

bique; Myanmar; Namibia; Nepal; Nigeria Pakistan; Peru; Philippines; Rwanda; Senegal; Sri Lanka; Tanzania; Thailand; Uganda; Vietnam; Zambia.

A.2 Additional Results

Table 15: Convergence Regressions for EETD \cap ASD/PLD (EETD, 1990–2018)

	Aggregate		Manufacturing		Agriculture	
	(1)	(2)	(3)	(4)	(5)	(6)
β	-0.005 (0.002)	-0.029 (0.010)	-0.001 (0.002)	-0.052 (0.009)	-0.001 (0.001)	-0.128 (0.035)
Number of observations			952			
Number of countries			34			
Units	Constant international prices from 2005					
Time fixed effects	Yes					
Country fixed effects	No	Yes	No	Yes	No	Yes

Standard errors clustered at country level are in parentheses.

Table 16: Geographic Robustness of Convergence Regressions (EETD, 1990–2018)

	Aggregate		Manufacturing		Agriculture	
	(1)	(2)	(3)	(4)	(5)	(6)
Sub-Saharan African countries excluded						
β	-0.010 (0.002)	-0.033 (0.011)	-0.005 (0.003)	-0.055 (0.023)	-0.003 (0.001)	-0.140 (0.033)
Observations			1,288			
Number of countries			46			
South and East Asian countries excluded						
β	-0.003 (0.001)	-0.055 (0.021)	0.003 (0.002)	-0.050 (0.013)	0.0001 (0.001)	-0.164 (0.036)
Observations			1,232			
Number of countries			44			
Latin American countries excluded						
β	-0.005 (0.001)	-0.036 (0.010)	-0.003 (0.002)	-0.043 (0.014)	-0.002 (0.001)	-0.126 (0.024)
Observations			1,540			
Number of countries			55			
Units	Constant international prices from 2005					
Time fixed effects	Yes					
Country fixed effects	No	Yes	No	Yes	No	Yes

Standard errors clustered at country level are in parentheses.

A.3 The Role of Different Productivity Measures

To assess the role of differences in data construction, we run the convergence regression (2) on the 64 countries of the EETD for both our productivity measure in constant international prices from 2005 and for Rodrik’s productivity measure in current USD prices. As before, we run the regressions both with and without country fixed effects. Different than before, we run the regressions both for our entire sample period 1990–2018 and for the subperiod 1995–2005. The decade 1995–2005 is informative because Rodrik (2013) paid particular attention to it, reflecting that the UNIDO data stop in 2005 and have the most balanced sample during 1995–2005.

Table 17 presents the results. The main message of the table is that the differences in methods for calculating comparable productivity measures is not crucial for the unconditional β -convergence results. In particular, while the point estimates using Rodrik’s data construction are somewhat larger than using our data construction for the entire 1990–2018 period, they are still small compared to his point estimates. Moreover, when focusing on the 1995–2005 subperiod, both β estimates for unconditional convergence are effectively zero. When considering conditional convergence, each data construction method implies significant convergence for both the overall time period and the 1995–2005 subperiod, though the estimates of conditional convergence are somewhat larger using the Rodrik construction.

We conclude that the starkly different results that we obtain for unconditional convergence compared to Rodrik are not due to the different methods for constructing comparable productivity levels.

Table 17: Convergence Regressions for Manufacturing – Constant International Prices vs. Current Prices in USD (EETD)

	(1)	(2)	(3)	(4)
β for 1990–2018	–0.003 (0.022)	–0.046 (0.014)	–0.006 (0.002)	–0.085 (0.018)
Observations	1,792			
β for 1995–2005	–0.001 (0.002)	–0.104 (0.035)	0.001 (0.004)	–0.260 (0.034)
Observations	640			
Number of countries	64			
Units	Constant 2005 int. prices		Current prices in USD	
Time fixed effects	Yes			
Country fixed effects	No	Yes	No	Yes

Standard errors clustered at country level are in parentheses.