UNCONDITIONAL CONVERGENCE IN MANUFACTURING*

Dani Rodrik

Unlike economies as a whole, manufacturing industries exhibit strong unconditional convergence in labor productivity. The article documents this at various levels of disaggregation for a large sample covering more than 100 countries over recent decades. The result is highly robust to changes in the sample and specification. The coefficient of unconditional convergence is estimated quite precisely and is large, at between 2–3% in most specifications and 2.9% a year in the baseline specification covering 118 countries. The article also finds substantial sigma-convergence at the two-digit level for a smaller sample of countries. Despite strong convergence within manufacturing, aggregate convergence fails due to the small share of manufacturing employment in low-income countries and the slow pace of industrialization. Because of data coverage, these findings should be as viewed as applying to the organized, formal parts of manufacturing. *JEL* Codes: O40, O14.

I. Introduction

Neoclassical growth theory establishes a presumption that countries with access to identical technologies should converge to a common income level. Countries that are poorer and have higher marginal productivity of capital should grow more rapidly in the transition to the long-run steady state. In an open global economy, access to foreign capital and foreign markets (which removes finance and market size as constraints) further strengthens the presumption of convergence.

However, empirical work has not been kind to this proposition. Selected developing countries, such as those in East Asia, have grown quickly. But when poor countries are taken as a whole, there is no systematic tendency for them to grow faster than rich ones, over any reasonably long time horizon for which we have data.[1] Whatever convergence one can find is conditional: It depends on policies, institutions, and other country-specific circumstances. The only clear-cut exceptions to the rule seem to be

*This is a substantially revised version of Rodrik (2011a). I am grateful to UNIDO for making the INDSTAT2 and INDSTAT4 databases available. I also thank Cynthia Balloch for research assistance; the Weatherhead Center for International Affairs and the Center for International Development at Harvard University for financial assistance; and Alberto Abadie, Daron Acemoglu, Jonathan Temple, three referees, and the editors for useful suggestions.

1. On convergence in the decade before the global financial crisis of 2008–2009, see Subramanian (2011, chap. 4), and Rodrik (2011b).

The Quarterly Journal of Economics (2013), 165–204. doi:10.1093/qje/qjs047.
Advance Access publication on November 18, 2012.

states/regions within a unified economy such as the United States (Barro and Sala-i-Martin 1991).[2]

If growth rates are characterized by conditional instead of unconditional convergence, economies will tend toward different levels of income in the long run. Lack of empirical support for (unconditional) convergence has led theory in the direction of models with endogenous technological change, which do not necessarily exhibit convergence, and to empirical work that focuses on identifying the conditioning variables that makes convergence feasible (see Acemoglu 2009 on theory, and Durlauf, Johnson, and Temple 2005 on empirical work).

In contrast to this large literature, I show in this article that unconditional convergence does exist, but it occurs in the modern parts of the economy rather than the economy as a whole. In particular, I document a highly robust tendency toward convergence in labor productivity in *manufacturing* activities, regardless of geography, policies, or other country-level influences. The coefficient of unconditional convergence (beta) is large—2.9% a year in my baseline specification that covers 118 countries—and estimated quite precisely, with more disaggregated specifications generally yielding somewhat higher estimates. A convergence rate of 2.9% implies that industries that are, say, a tenth of the way to the technology frontier (roughly the bottom 20% of the industries in the sample) experience a convergence boost in their labor productivity growth of 6.7 percentage points per annum ($0.029 \times \ln(10)$).

Figures I, II, and III illustrate the central result of this article and place it in proper perspective. Each figure shows the relationship between labor productivity in some base period (on the vertical axis) and its growth rate over the subsequent decade, controlling for period-specific influences. Figure I, where each dot stands for a particular country in a specific decade, presents a typical nonconvergence result for country-level productivity. There is no systematic tendency for countries that start with lower productivity (measured here by gross domestic product [GDP] per worker) to grow more rapidly.

2. Some studies also find unconditional convergence among the richer Organisation for Economic Co-operation and Development (OECD) countries, but it is difficult to know what to make of this result in light of the obvious sample selection bias (Baumol 1986; DeLong 1988).

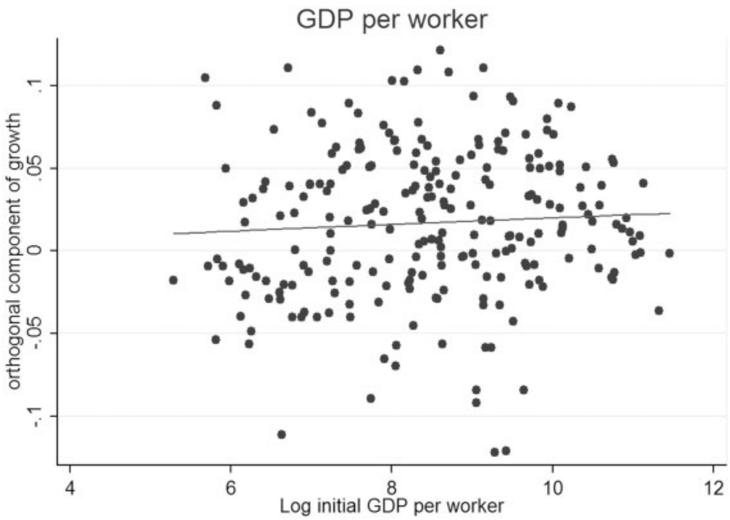

FIGURE I

Lack of Convergence in Economy-wide Labor Productivity

Variable on the vertical axis is growth of GDP per worker over four separate decades (1965–1975, 1975–1985, 1985–1995, 1995–2005), controlling for decadal fixed effects. Source of data: PWT 7.0. Sample is restricted to countries included in the manufacturing convergence regressions.

Figure II depicts the analogous relationship for individual manufacturing industries. Each dot on this scatter plot represents a two-digit industry in a specific country. (Illustrative industries: food and beverages, chemicals and chemical products, motor vehicles.) Each country enters Figure II with multiple industries (but over a single time horizon, with the most recent decade for which it has data). Period and industry-specific influences are controlled using industry, decade, and industry $\times$ decade fixed effects. Because there are no controls for country-level determinants, Figure II represents a test of unconditional convergence similar in spirit to Figure I. (The need for period and industry fixed effects are motivated subsequently.) The negative and highly significant slope is unmistakable, illustrating the central conclusion: Manufacturing exhibits a strong tendency for unconditional convergence. Industries that start at lower levels of labor productivity experience more rapid growth in labor

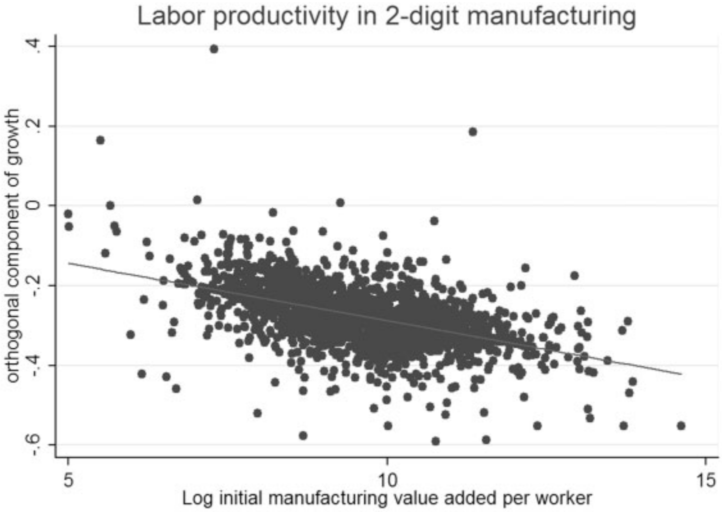

FIGURE II

Unconditional Unconditionce in 2-digit Manufacturing Sectors

Variable on the vertical axis is the growth of value added per worker in 2-digit manufacturing industries, controlling for period, industry, and period × industry fixed effects, where for each country the latest decade over which data are available is included. Source of data: INDSTAT2. For further details on data and methods, see text.

productivity. As I show later, when country fixed effects are included the slope becomes even steeper. Conditional convergence is more rapid than unconditional convergence. But what is striking in Figure II is the evident strength of convergence in the data even in the absence of such controls. Figure III shows the same convergence result for manufacturing in aggregate, with labels identifying each country.[3]

My focus in this article is on what it is called beta-convergence, where *beta* refers to the slope of the relationship in Figures I–III. I also present evidence for significant

3. There are some apparent anomalies in Figure III arising from compositional differences across countries. For example, Suriname (SUR) shows up as the country with the highest labor productivity. This is due to the fact that industry in this country consists almost entirely of alumina and aluminum production, which is highly capital- and energy-intensive. Such compositional issues provide an important rationale for working with more disaggregated data.

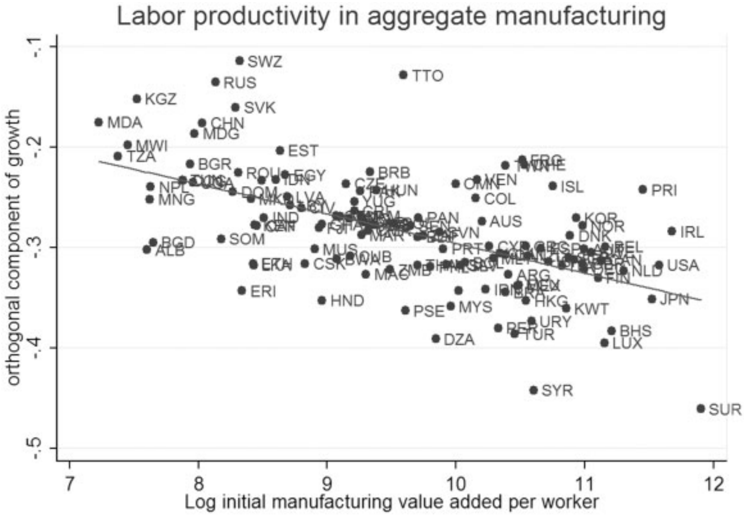

FIGURE III

Unconditional Convergence in Aggregate Manufacturing

Variable on the vertical axis is the growth of value added per worker in aggregate manufacturing controlling for period fixed effects, where for each country the latest decade over which data are available is included. Source of data: INDSTAT2. For further details on data and methods, see text.

sigma-convergence, referring to convergence in productivity *levels*, at least for recent time periods. Even if beta-convergence holds, countries may fail to converge in levels as long as random shocks to the growth process are relatively large.[4] In the present sample of industries, beta- and sigma-convergence go together. I find that productivity dispersion was sharply reduced across countries during 1995–2005 in the majority of two-digit manufacturing industries. In manufacturing taken as a whole, sigma shrunk by 10 log-points during this period in the sample of 63 countries for which comparison is possible—a reduction of the order of 10%.

4. This is true especially if the shocks act in an offsetting manner, but sigma-convergence may fail even if the shocks are purely idiosyncratic (i.e., uncorrelated with incomes per capita). Young, Higgins, and Levy (2008) discuss the difference between beta- and sigma-convergence and document the absence of sigma-convergence for U.S. counties.

I note that my data come from the United Nations Industrial Development Organization's (UNIDO) industrial statistics database, which is derived largely from industrial surveys. Because microenterprises and informal firms are often excluded from such surveys, we cannot be certain that the results are universally valid across all types of manufacturing activities. In the absence of more complete coverage of manufacturing, these findings on convergence should be as viewed as applying to the organized, formal parts of manufacturing.

To my knowledge, this is the first article to demonstrate unconditional convergence in industry for a wide range of countries and for detailed manufacturing industries. There does not seem to be any work that has looked at highly disaggregated data for manufacturing or at the manufacturing experience of countries beyond Organisation for Economic Co-operation and Development (OECD) and U.S. states (Bernard and Jones 1996a. 1996b; see also Sørensen 2001). However, two recent related studies deserve mention. In unpublished work, Hwang (2007) has documented that there is a tendency for unconditional convergence in export unit values in highly disaggregated product lines. Once a country begins to export something, it travels up the value chain in that product regardless of domestic policies or institutions.[5] Hwang shows that the lower the average unit values of a country's manufactured exports, the faster the country's subsequent growth, unconditionally. This article differs from Hwang in that it focuses on output rather than exports and directly on productivity (rather than unit values). Convergence seems to kick in manufacturing regardless of whether production is exported.[6] In addition, a recent paper by Bénétrix, O'Rourke, and Williamson (2012) has documented convergence in industrial output since the late nineteenth century, with the strongest results obtaining for the 1890–1972 period. Bénétrix and colleagues focus on growth rates of aggregate industrial

5. Hwang demonstrates his result for both 10-digit U.S. HS import statistics and 4-digit SITC world trade statistics. The first classification contains thousands of separate product lines.

6. Also related is a paper by Levchenko and Zhang (2011), which estimates model-based relative productivity trends for 19 manufacturing industries from the 1960s through the 2000 and show that there has been steady convergence across countries.

output rather than on disaggregated industrial productivity, which is the focus here.[7]

These results naturally raise the question of why convergence within manufacturing fails to aggregate up to the economy as a whole. Another contribution of this article is to reconcile unconditional convergence in manufacturing with its apparent absence for GDP per worker. I analyze this question by combining UNIDO data with Penn World Tables to derive implied labor productivity estimates for the rest of the economy. With these data, I document three additional features of the data. First, nonmanufacturing does not exhibit convergence. Second, manufacturing's impact on aggregate convergence is curtailed by its small size, especially in the poorer countries. Third, the growth boost from reallocation—the shift of labor to more productive manufacturing—is not sufficiently and systematically greater in poorer economies. Taken together, these three facts account for the absence of aggregate convergence.

The analysis highlights the role of structural factors, in particular the slow (and sometimes perverse) movement of resources across economic activities with different convergence characteristics. The trouble from a convergence standpoint is that economic activities that are good at absorbing advanced technologies are not necessarily good at absorbing labor. As a result, too much of an economy's resources can get stuck in the "wrong" sectors—those that are not on the escalator up. When firms that are part of international production networks or otherwise benefit from globalization employ little labor, the gains remain limited. Even worse, intersectoral labor flows can be perverse with the consequence that convergence within the "advanced" sectors is accompanied by divergence on the part of the economy as a whole. McMillan and Rodrik (2011) illustrate some of these perverse outcomes using the experience of specific countries.

The article proceeds as follows. Section II describes the data and methods used for the estimation. Section III presents the basic results and various robustness checks. Section IV considers the reasons for convergence failing to aggregate up to the level of the entire economy. Section V provides concluding remarks.

7. The first version of this article (Rodrik 2011a) was completed before I became aware of Bénétrix, O'Rourke, and Williamson's (2012) work on the subject.

II. Data and Methods

II.A. Date Source and Description

I use data from UNIDO's INDSTAT2 and INDSTAT4 databases, which provide industrial statistics for a wide range of countries at different levels of disaggregation (UNIDO 2011, 2012). My baseline sample is based on INDSTAT2, which has good coverage of countries at the ISIC two-digit level going back to the 1960s. These statistics cover value added and employment, among others, for up to 23 manufacturing industries per country, allowing me to calculate labor productivity (value added per employee) and its growth at that level of disaggregation. INDSTAT4 provides more disaggregated data at the four-digit level for up to 127 industries, but covers fewer countries and is spotty for earlier years, making it impractical to work with it for periods that extend before 1990. An earlier version of this article reported substantially similar results using INDSTAT4 (Rodrik 2011a). Here, I rely mostly on INDSTAT2, which has the advantage of allowing me to increase the country coverage as well as present results for before 1990.[8] Results using INDSTAT4 serve to provide robustness checks.

Even with INDSTAT2 I am hampered by the fact that few countries have long time series, with developing countries posing a particular challenge. The longer the time span we choose for the convergence horizon, the smaller the number of countries that can be included. This makes any horizon longer than a decade impractical. Accordingly, the empirical analysis is based on decadal growth. Moreover, there is limited overlap across countries for any particular decade, so that choosing a fixed time span—say, 1995 to 2005—reduces the sample considerably, by half or more.

In light of these limitations, my baseline consists of a pooled sample that combines the latest 10-year period for each country with data. The advantage of pooling is that it maximizes the number of countries that can be included, yielding a sample of 118 countries in all and allowing good coverage of poorer economies.

8. INDSTAT2 was released after Rodrik (2011a) was completed. I thank a referee for bringing it to my attention.

Because each country enters with around 20 industries, the total number of observations in the baseline specification is greater than 2,000. I supplement this sample with two others to check for robustness and carry out further analyses:

(1) A panel, which stacks the four decades 1965–75, 1975–85, 1985–95, and 1995–2005 for all countries with data for at least one of those decades. The panel sample covers a relatively large number of countries (99 in all). Its disadvantage is that it is highly unbalanced, with developed countries having much better coverage than developing ones. For example, Japan enters the panel with 71 industries, whereas Algeria has merely 12.

(2) A pure cross-section for 1995–2005. This sample has the fewest countries covered (58), for reasons already explained. I rely on the cross-section sample when I require a consistent comparison across countries over a fixed time period (as when I analyze sigma-convergence or examine the failure of aggregation).

Some countries in INDSTAT2 have data for aggregate manufacturing, but lack disaggregated data. Thus the samples become slightly larger than what was just noted when I analyze manufacturing as a whole.

As mentioned in Section I, UNIDO's data come from industrial surveys whose coverage differs across countries. Data for developed countries refer for the most part to "all establishments." But in developing countries, enterprises with fewer than 5 or 10 employees are often not included. For this reason, the convergence results that follow should be read as applying to the more formal, organized parts of manufacturing and not to micro-enterprises or informal firms. An appendix (available on request) provides a summary of data coverage for each country included in the regressions.

An important problem with the data is that INDSTAT provides figures for value added in nominal U.S. dollars. What I need is a measure of growth in labor productivity in real terms. I can recover the convergence parameter under certain assumptions about the process followed by U.S. dollar inflation rate for disaggregated manufacturing industries. In what follows, I explain my approach in greater detail.

II.B. Empirical Specification

Dividing nominal US$ value added by employment, I calculate nominal labor productivity for each industry v_{ijt}, where i denotes the industry, j the country, and t the time period. The rate of growth of labor productivity in real terms, $\hat{y}_{ijt}$, is given in turn by $\hat{y}_{ijt} = \hat{v}_{ijt} - \pi_{ijt}$, where π_{ijt} is the increase in the industry-level deflator in dollar terms and a hat over a variable denotes percent changes.

I assume (real) labor productivity growth in each industry is a function of both country-specific conditions and a convergence effect. The latter, in turn, is proportional to the gap between each industry's initial productivity and its frontier technology, represented by v^*_{it}. Hence:

$$\hat{y}_{ijt} = \beta(\ln v^*_{it} - \ln v_{ijt}) + D_j,$$

where D_j is a dummy variable that stands in for all time- and industry-invariant country-specific factors. The convergence coefficient we are interested in estimating is β. Note that if $\ln v_{ijt}$ is measured with error, this specification potentially introduces a bias toward *over*estimating the rate of convergence, since such an error weakens the link between initial productivity and final productivity. This is a common problem in the empirical literature on convergence (Temple 1998).

The last step is to specify a process for prices. I assume a common global (U.S. dollar) inflation rate for each individual industry up to an idiosyncratic (random) error term, such that $\pi_{ijt} = \pi_{it} + \varepsilon_{ijt}$. This is a reasonable assumption because manufactures are tradable and face common world prices. Of course, in practice there are many reasons domestic prices may diverge from world prices, even in tradables. Transport costs or trade policies such as import tariffs and export subsidies drive wedges between domestic and foreign prices. But these wedges introduce differences in price levels, not inflation rates. Equivalently, I assume that dollar inflation rates are not systematically correlated with an industry's distance from the technological frontier. My results do not overstate convergence as long as dollar price inflation is not systematically higher in industries that are furthest away from the technological frontier. I postpone further discussion of these issues to the next section, where I discuss potential complications arising from the absence of price information more extensively.

This allows me to express the growth of nominal labor productivity as follows:

$$\hat{v}_{ijt} = -\beta(\ln v_{ijt} - \ln v^{*}_{it}) + \pi_{it} + D_j + \varepsilon_{ijt}. \tag{1}$$

I assume ε_{ijt} is uncorrelated with other explanatory variables and captures all other idiosyncratic influences on measured labor productivity. Rearranging terms, I now have the final estimating equation:

$$\hat{v}_{ijt} = -\beta \ln v_{ijt} + (\pi_{it} + \beta \ln v^{*}_{it}) + D_j + \varepsilon_{ijt}. \tag{2}$$

This can be expressed equivalently as

$$\hat{v}_{ijt} = -\beta \ln v_{ijt} + D_{it} + D_j + \varepsilon_{ijt}, \tag{3}$$

where D_{it} stands for $(\pi_{it} + \beta \ln v^{*}_{it})$. Hence, I can regress the growth of labor productivity in nominal U.S. dollar terms on the initial level of labor productivity, a set of industry $\times$ time period fixed effects (D_{it}), and country fixed effects (D_j). In the regressions that follow, I include separate industry and period dummies to soak up any additional confounding residual systematic variation.

It is also possible to run this regression over a single time period as a pure cross-section. In this case, the industry $\times$ time period fixed effects are reduced to industry fixed effects:

$$\hat{v}_{ij} = -\beta \ln v_{ij} + D_i + D_j + \varepsilon_{ij}. \tag{3$'$}$$

As specified, the estimate of β will be a measure of *conditional* convergence, since country-specific conditions are explicitly controlled for by the country fixed effects. A test of *unconditional* convergence consists of dropping these country dummies and checking whether the estimated coefficient $-\beta$ remains negative and statistically significant.

III. Empirical Results

III.A. Basic Results

Table I shows the results for the baseline specification along with its variants. The dependent variable in each case is the (compound annual) growth rate of labor productivity for two-digit manufacturing industries. The regressors are the log of initial

TABLE I

BASELINE SPECIFICATION AND VARIATIONS (DEPENDENT VARIABLE: GROWTH OF LABOR PRODUCTIVITY IN TWO-DIGIT MANUFACTURING INDUSTRIES OVER A DECADE)

	Baseline		Baseline, only post-1990 data		Cross-section, 1995–2005		Panel			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Log initial productivity	−0.029***	−0.054***	−0.029***	−0.057***	−0.024***	−0.059***	−0.018***	−0.060***	−0.019***	−0.055***
	(0.004)	(0.006)	(0.004)	(0.006)	(0.004)	(0.010)	(0.003)	(0.005)	(0.006)	(0.007)
Log initial productivity, 1975–85									0.000	−0.009
									(0.006)	(0.006)
Log initial productivity, 1985–95									0.016*	0.008
									(0.009)	(0.007)
Log initial productivity, 1995–2005									−0.006	−0.011
									(0.007)	(0.009)
Country fixed effects	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Industry fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Period fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Period × industry fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of countries	118	118	104	104	58	58	99	99	99	99
Number of observations	2,122	2,122	1,861	1,861	955	955	3,848	3,848	3,848	3,848

Notes. The baseline specification uses all countries with the requisite data, using only the latest 10-year period for each. The panel specification pools data for 1965–75, 1975–85, 1985–95, and 1995–2005 for all countries with requisite data. Standard errors are clustered at country level.

$^{***}p < .01$, $^{**}p < .05$, $^{*}p < .10$.

labor productivity and a host of fixed effects, depending on the specification. Each regression is run first without and then with country dummies. As explained previously, these two specifications yield the unconditional and conditional convergence coefficients, respectively. Standard errors are clustered at the country level in all specifications.

Column (1) is the baseline result, corresponding to the scatter plot displayed in Figure II. The estimated convergence coefficient of 2.9% is highly significant, with a *t*-statistic of 6.95. Recall that the baseline sample pools different decades for each country. There is little evidence of parameter heterogeneity across different time periods, however. I get quite similar estimates when I drop pre-1990 observations from the baseline specification (column (3)) or when I run a pure cross-section for the more recent period 1995–2005 (column (5)).[9] I take note of the fact that the 1995–2005 results do not diverge much from the baseline as I have to focus on the cross-section for some of the subsequent analysis.

The panel specification, which combines the four decades 1965–75, 1975–85, 1985–95, and 1995–2005, yields a somewhat smaller β of 1.8%, but this estimate is still highly significant with a *t*-statistic of 6.74 (column (7)). In column (9) I test explicitly for differences in the convergence coefficient across time periods by interacting base labor productivity with decade dummies. The excluded decade is 1965–75. The results suggest a lower convergence estimate for 1985–95 but otherwise no strongly discernible trends over time. In particular, there is no evidence of stronger convergence in more recent decades. This is somewhat surprising because one might have expected globalization and the spread of global production networks to greatly facilitate technological dissemination and therefore catch-up. The result suggests convergence is an intrinsic property of manufacturing industries, and one that is not driven by the ups and downs of global economic integration.

Each specification in Table I is paired with its *conditional* variant, which includes country fixed effects. The estimated convergence coefficients always increase in size, typically doubling—or tripling, in the panel specification—when country dummies

9. The cross-section results for the latest decade for which I have data, 1997–2007, are very similar to those for 1995–2005 (not shown). I present the 1995–2005 results because of the somewhat larger country coverage.

are included. This is in line with the conditional convergence results in the literature. As noted recently by Barro (2012), there are reasons to think growth regressions with country fixed effects yield upwardly biased estimates of the convergence rate when the time horizon is short. This is due to the so-called Hurwicz-Nickell bias: the estimate of the coefficient on the lagged dependent variable is biased downwards in the presence of a fixed effect (Hurwicz 1950; Nickell 1981). The bias tends to zero as the time span gets large, but can be large in short panels.[10] For my purposes, because I am mainly interested in unconditional convergence (and hence the specifications without country fixed effects), it suffices to treat the conditional convergence estimates as upper bounds. They still provide a useful reference point for gauging the magnitude of the unconditional convergence rates.

It is to be expected that country-specific conditions that are correlated with initial productivity—policies, institutions, geography—play a role in determining the speed of catch-up. So even in the absence of the Hurwicz-Nickell bias, it is natural that the beta coefficient would become larger when fixed effects are introduced. What is surprising is the evidence for systematic and rapid productivity convergence in individual manufacturing industries when these country-specific conditions are not controlled for through country dummies. Once again, there is no evidence of significant differences in convergence rates across time periods. In particular, the speed of (conditional) convergence does not appear to have increased in recent decades (column (10)).

In Table II, I compare the baseline with results at different levels of aggregation. The table shows specifications that are both more disaggregated (three- and four-digit, using INDSTAT4 data) and more aggregated (manufacturing as a whole). The results in column (7) correspond to Figure III. Each regression is run in two versions, one with the largest possible sample and another one that restricts the sample to a common list of countries with the requisite data at all levels of aggregation. The latter facilitates direct comparison across different levels of aggregation. The estimated βs are statistically significant, typically at

10. In principle, the same bias exists when industry fixed effects are included too, as in my baseline specification. But the dimensionality is much greater in this case since the de-meaning of the dependent variable takes place over both time periods and across countries. In any case, I also provide pure cross-section estimates for each industry separately, using no fixed effects.

TABLE II

COMPARISON OF RESULTS AT DIFFERENT LEVELS OF AGGREGATION (DEPENDENT VARIABLE: GROWTH OF LABOR PRODUCTIVITY IN MANUFACTURING INDUSTRIES OVER A DECADE)

	Baseline (ISIC 2-digit)		ISIC 3-digit		ISIC 4-digit		All manufacturing	
	(1) All	(2) Common sample	(3) All	(4) Common sample	(5) All	(6) Common sample	(7) All	(8) Common sample
Log initial productivity	−0.029***	−0.028***	−0.029***	−0.030***	−0.031***	−0.031***	−0.030***	−0.023**
	(0.004)	(0.008)	(0.006)	(0.006)	(0.004)	(0.004)	(0.006)	(0.009)
Country fixed effects	No	No	No	No	No	No	No	—
Industry fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	No	—
Period fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Period × industry fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	No	—
Number of countries	118	38	39	38	39	38	123	38
Number of observations	2,122	812	704	703	2,871	2,869	123	38

Notes. Standard errors are clustered at country level, except for "all manufacturing," for which robust standard errors are reported. Common sample refers to the set of countries that have data at all levels of aggregation showed in this table, for the same period.

$^{***}p < .01$, $^{**}p < .05$, $^{*}p < .10$.

the 99% confidence level, at all levels of aggregation. The estimates at the three- and four-digit level are remarkably similar to the baseline estimates, as are the estimates for manufacturing as a whole. Looking across the regressions with a common sample, one can see some indication that the estimated coefficient increases with the level of disaggregation—from 2.3% for manufacturing as a whole to 3.1% at the four-digit level. But these differences are not statistically significant. The similarity in the magnitude of the estimates suggests that the failure of unconditional convergence to aggregate up to the economy as a whole is not a result of compositional issues within manufacturing. I make use of this result later in the article.

III.B. Robustness

Table III presents a battery of additional robustness tests around the baseline specification. Panel A prunes the baseline sample in a number of ways to ensure that countries/industries with specific characteristics are not driving the results. I exclude in turn (1) countries with fewer than 10 industries; (2) observations that correspond to the highest and lowest 10% values for growth; (3) observations in the top and bottom half, respectively, of the sample in terms of labor productivity; (4) former socialist countries; and (5) OECD countries.

Remarkably, the convergence estimates remains highly significant across all these runs. They vary from a low of 1.1% when growth observations at the top and bottom deciles are excluded (column (3)) to a high of 6.0% for the low-productivity half of the sample (column (4)). Note that the result becomes, if anything, stronger when OECD countries are excluded (column (7)), with the convergence coefficient rising from 2.9% to 3.8%. This is significant, because unconditional convergence has never been documented outside the OECD.

Panel B of Table III carries out different types of robustness tests, including weighting observations by value added, instrumenting initial labor productivity by lagged productivity (a check against measurement error), recalculating growth rates by estimating a log-linear trend using all 10 annual observations (instead of just endpoints),[11] and clustering standard errors by

11. This guards against introducing a spurious bias that may arise from lagged labor productivity appearing directly on both sides of the regression equation, with opposite signs.

TABLE III

ROBUSTNESS TESTS (DEPENDENT VARIABLE: GROWTH RATE OF PRODUCTIVITY OVER RELEVANT PERIOD)

Panel A							
	(1) Baseline	(2) Excluding countries with fewer than 10 industries included	(3) Excluding observations with the highest and lowest 10% values for growth	(4) Subsample with labor productivity in the bottom half of the sample	(5) Subsample with labor productivity in the top half of the sample	(6) Excluding former socialist countries	(7) Excluding OECD countries
Log initial productivity	−0.029***	−0.027***	−0.011***	−0.060***	−0.043***	−0.025***	−0.038***
	(0.004)	(0.004)	(0.003)	(0.006)	(0.005)	(0.004)	(0.006)
Country fixed effects	no	no	no	no	no	no	no
Industry fixed effects	yes	yes	yes	yes	yes	yes	yes
Period fixed effects	yes	yes	yes	yes	yes	yes	yes
Period × industry fixed effects	yes	yes	yes	yes	yes	yes	yes
Number of countries	118	109	116	95	108	112	91
Number of observations	2,122	2,058	1,697	1,061	1,061	1,997	1,542

TABLE III
(CONTINUED)

Panel B

	(8) Weighted regression results: ln (value added) weights	(9) IV results with 5-year lagged initial productivity as instrument	(10) Rescaled labor productivity growth, adjusted for real exchange rate appreciation	(11) Standard errors clustered by industry	(12) Country fixed effects with no industry controls	(13) Growth rates calculated from log-linear trend using 10 annual observations	(14) Testing for nonlinearity by labor productivity quartiles
Log initial productivity	−0.030***	−0.024***	−0.030***	−0.029***	−0.021***	−0.017***	
	(0.004)	(0.005)	(0.004)	(0.003)	(0.003)	(0.006)	
Log initial productivity, lowest quartile							−0.034***
							(0.007)
Log initial productivity, second quartile							−0.035***
							(0.006)
Log initial productivity, third quartile							−0.034***
							(0.006)
Log initial productivity, fourth quartile							−0.032***
							(0.005)
Country fixed effects	no	no	no	no	yes	no	no
Industry fixed effects	yes	yes	yes	yes	no	yes	yes
Period fixed effects	yes	yes	yes	yes	yes	yes	yes
Period × industry fixed effects	yes	yes	yes	yes	no	yes	yes
Number of countries	118	104	102	23	118	107	118
Number of observations	2,122	1,615	1,754	2,122	2,122	1,774	2,122

Notes. Standard errors are clustered at country level, unless otherwise noted.
$^{***}p < .01$, $^{**}p < .05$, $^{*}p < .10$.

industry rather than country. In all these runs, the results remain highly significant. Note in particular that the estimated coefficient remains virtually unchanged when I weight industries by size (column (8)).[12] Hence, the convergence result is not driven by the experience of relatively small industries. Instrumenting for initial productivity reduces the estimated β marginally to 2.4% (column (9)).

The last column of Table III tests directly for nonlinearity of β by allowing the estimated coefficient to vary by labor productivity quartile. The results do not suggest any nonlinearity—at least at the two-digit level (column (14)). There is greater evidence of nonlinearity in the four-digit data (INDSTAT4), as I reported in the earlier version of this article (Rodrik 2011a).[13]

Another way to examine the data is to scrutinize the convergence evidence on an industry-by-industry basis. This exercise not only serves as a robustness test but is interesting in its own right. I begin with a few scatter plots for individual industries. Figure IV shows convergence plots for the four largest industries in the sample that also have good country coverage: ISIC 15 (food and beverages), 24 (chemicals and chemical products), 29 (machinery and equipment, n.e.c.), and 34 (motor vehicles). (In keeping with the baseline specification, each country enters these scatterplots with the latest period for which it has data. I continue to control for time-specific inflation trends by including period dummies.) Note that these plots use data from just the specified industries. As the negative slopes indicate, countries that started further behind tended to experience more rapid productivity growth in all four industries. The relationship is statistically significant at the 99% level in all four cases, with estimated coefficients ranging from 2.2% (ISIC 34) to 2.8% (ISIC 15).

The individual convergence coefficients estimated on an industry-by-industry basis for each of our two-digit industries are shown in Table IV. I regress, separately for each industry, the growth rate of an industry's labor productivity against its

12. Using employment weights produces nearly identical results.

13. In Rodrik (2011a) I also reported some evidence that industries differ in their β coefficients. These differences are much less evident in the two-digit data I use here. The only finding of note is that textiles and clothing show a mildly lower convergence coefficient. These results, not shown here, are available on request.

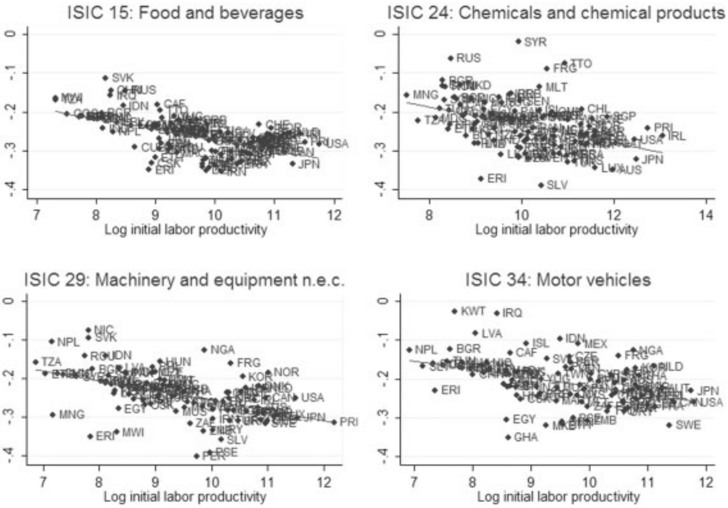

FIGURE IV

Convergence in Specific Industries

Vertical axis represents relative growth rate of labor productivity, controlling for period fixed effects. Each country enters with most recent decade for which data are available.

initial level across all countries in the sample that have the requisite data for the 1995–2005 period:

$$\hat{v}_{ij} = -\beta_{\mathrm{i}} \ln v_{ij} + \varepsilon_{ij} \quad \text{for} \quad i = 1, \ldots, \mathrm{I}.$$

This entails running as many regressions (23) as I have manufacturing industries. This is the direct analogue of running cross-country growth regressions. Note that these specifications do not contain industry, period, or any other fixed effects. Many of these regressions cover 30–40 countries, so one should not be too demanding in terms of statistical significance for industry-specific estimates. Nevertheless, the results are quite strong. Among the 23 industries, I find that 18 have statistically significant convergence coefficients, ranging from a low of 1.4% (textiles) to a high of 5.5% (office, accounting, and computing machinery). All but one of the industries exhibit unconditional convergence, with wearing apparel the sole exception.

TABLE IV

BETA-CONVERGENCE COEFFICIENTS BY INDUSTRY, 1995–2005 REGRESSIONS

ISIC code	Sector	Number of countries	Beta coefficient	Significance
D	Total manufacturing	63	−.020	***
15	Food and beverages	43	−.026	***
16	Tobacco products	31	−.020	**
17	Textiles	54	−.014	*
18	Wearing apparel, fur	34	.005	n.s.
19	Leather, leather products, and footwear	34	−.001	n.s.
20	Wood products (excl. furniture)	53	−.033	***
21	Paper and paper products	52	−.024	***
22	Printing and publishing	53	−.021	***
23	Coke, refined petroleum products, nuclear fuel	38	−.026	**
24	Chemicals and chemical products	52	−.010	n.s.
25	Rubber and plastics products	53	−.019	***
26	Nonmetallic mineral products	56	−.032	***
27	Basic metals	50	−.046	***
28	Fabricated metal products	53	−.022	**
29	Machinery and equipment n.e.c.	37	−.011	n.s.
30	Office, accounting, and computing machinery	23	−.055	***
31	Electrical machinery and apparatus	32	−.023	***
32	Radio, television, and communication equipment	26	−.043	***
33	Medical, precision, and optical instruments	45	−.031	***
34	Motor vehicles, trailers, semi-trailers	35	−.019	*
35	Other transport equipment	27	−.030	**
36	Furniture; manufacturing n.e.c.	52	−.031	**
37	Recycling	22	−.019	n.s.

Notes. These coefficients are obtained by running pure cross-section regressions for each industry separately.
***$p < .01$, **$p < .05$, *$p < .10$, n.s. = not significant.

For an even more disaggregated analysis, I turn to INDSTAT4, which breaks down manufacturing into 127 separate subsectors. This allows a more fine-tuned scrutiny of the convergence experience of individual industries, albeit for a reduced country sample. As I have shown in Rodrik (2011a, table 2), four-digit data yield highly significant global convergence coefficients when all four-digit industries are pooled together, even for periods that are as short as five years. Here I focus on the results of individual industry regressions for 2000–2005, because this maximizes the number of countries (and hence observations). Once again, these are simple cross-section regressions for each industry separately, with no fixed effects.

Figure V summarizes the results by showing the distribution of estimated coefficients on initial productivity across the 127 industries, normalized by their standard errors. The vast majority of the estimates are negative. Though not all of them are statistically significant, a surprising number of the negative ones are. Specifically, 79 (out of 127) of the industry regressions yield negative and statistically significant coefficients (at the 95% level or higher). By contrast, none of the (few) positive coefficients are statistically significant.

There is a statistical sense in which the global estimate, which pools across industries and controls for industry fixed effects, represents a weighted average of these industry-by-industry convergence estimates. Specifically, consider a given cross-section (such as 2000–2005 in the foregoing). Omitting time subscripts, the relationship between estimates β and β_i is given by:

$$\beta = \sum_{i=1}^{I} \beta_{\mathrm{i}} \left(\frac{var\left(ln\, v_{ij} | \mathcal{J} = \mathrm{i}\right) \Pr(\mathcal{J} = i)}{\sum_{l=1}^{I} var\left(ln\, v_{ij} | \mathcal{J} = l\right) \Pr(\mathcal{J} = l)} \right),$$

where $\mathcal{J}$ denotes an industry-identifying variable. Thus, the weight that each industry estimate gets in the global estimate depends on the variance of the independent variable $\ln v_{ij}$ within the industry as well as the relative number of occurrences of that industry in the global sample.[14]

14. The derivation of this expression, for which I am indebted to Alberto Abadie, is available on request.

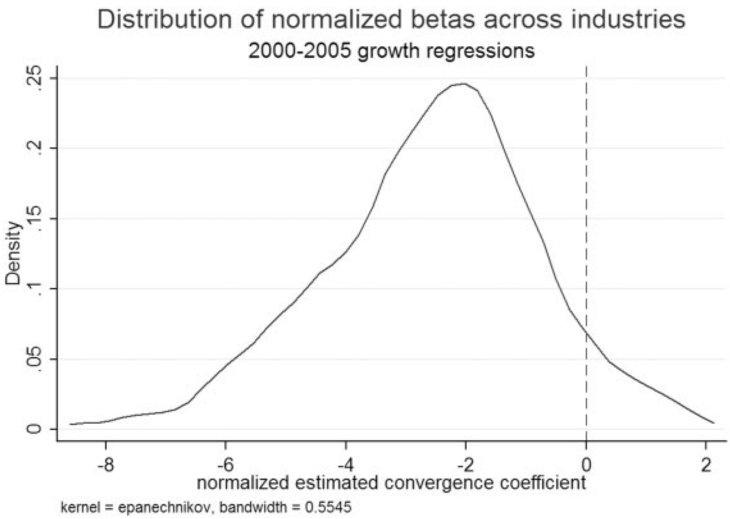

FIGURE V

Distribution of Convergence Coefficients, by 4-digit Industry (Normalized by Standard Error)

The figure is based on convergence coefficients generated from 127 industry convergence regressions across 4-digit manufacturing industries. See text.

III.C. Complications Arising from Price Effects

One possible concern in interpreting these results is that my assumption of a common value added deflator (in dollars) for each industry, regardless of the country where it is located, may be introducing a bias to the estimation. I justified this assumption previously by arguing that the manufacturing industries in question are tradable, and hence face common world prices. Of course, I do not expect the law of one price to obtain perfectly, even for homogeneous or standardized goods. Tariffs, subsidies, nontariff barriers, and transport costs often drive a wedge between domestic and world prices. Such wedges may well be larger in the poorer countries. My estimation strategy can accommodate such deviations provided they do not *vary over time* in a way that is correlated systematically with distance from the technological frontier.

Suppose, for example, the ad valorem equivalent of trade costs is τ_{ijt}. Then $\pi_{ijt} = \pi^*_{it} + (1 + \tau_{ijt})^{-1} d\tau_{ijt}$, where π^*_{it} is the percent change in world prices. All I need for unbiased estimation is for $(1+\tau_{ijt})^{-1}d\tau_{ijt}$ to be uncorrelated with initial labor productivity $\ln v_{ijt}$, conditional on a set of industry, period, and industry $\times$ period fixed effects. It is hard to think of strong reasons a priori as to why there should be such a pattern in trade cost changes.

One potential source of bias is systematic changes in real exchange rates. In principle, across-the-board increases in domestic costs such as wages should be offset, on average, by depreciation of the home currency, leaving dollar values generally unchanged. But in countries that experience sustained movements in their real exchange rate, trends in value added expressed in U.S. dollars will be misleading with regard to productivity in individual manufacturing industries. The worst case, from the perspective of the present article, would be if the low-income countries that house a preponderant share of low-productivity industries were the ones to experience real exchange rate appreciations—a rise in domestic costs not compensated by currency depreciation. This would lead to an upward bias in my convergence estimates.

To check against this possibility, Table III provides a version of the convergence regression that explicitly "corrects" for real exchange rate changes. I rescaled the growth in value added per worker by deflating it with (one plus) the rate of appreciation of the country's real exchange rate.[15] This reduces—across the board—the measured productivity growth rates of industries in countries which have experienced real appreciations, while raising them in countries with real depreciations. I rerun the baseline specification using these adjusted values for the dependent variable (Table III, column (10)).

If observed convergence were due to real appreciation in the poorer countries, the resulting estimates would be substantially lower. In fact, the estimated β is actually slightly higher (.030) and statistically equally significant. (Note, however, that the sample size is somewhat reduced as the lack of price data for some of the countries prevents me from computing real exchange

15. These are conventional bilateral real exchange rates vis-à-vis the United States. Domestic inflation rates have been calculated using producer-price indices where possible, substituting the Consumer Price Index where the Producer Price Index is not available. The source for the data on exchange rates and price indices is the International Monetary Fund's International Financial Statistics. A few countries had to be dropped because of lack of price data.

rate changes for them.) The bottom line is that there is no evidence that real exchange rate movements have distorted my basic findings.

Another source of possible bias arises from compositional changes. Even four-digit industries are a mix of different activities, and what appears as an increase in dollar values may be in reality a shift toward higher value-added activities within the same industry. Equivalently, producers may be moving toward higher-quality varieties, so that what looks like productivity growth is really quality upgrading (Schott 2004; Hwang 2007). It is possible that industries that start further away from the frontier experience such shifts more rapidly. In the presence of price data, I would have been able to capture such changes through their effect on industry-specific prices.

Such biases are potentially of concern, but my results provide some comfort that they are not very important in practice. As seen in Table II, the estimated speed of convergence changes very little across two-, three-, and four-digit levels of aggregation. If anything, the estimate increases in value as I disaggregate further. If the observed productivity growth was due to a shift toward higher quality products, the opposite result would hold.

Furthermore, even if the compositional biases were quantitatively significant, they would not detract greatly from the convergence result that is my focus. Moving into more sophisticated, higher value-added products and increasing physical productivity are both ways of raising the available returns to labor. They are both channels for income convergence. To the extent that quality upgrading takes place generally, and it does so more rapidly in the poorer countries, it is simply another manifestation of the productive convergence I am interested in documenting.

A final interpretational difficulty relates to the possible role of entry barriers in accounting for my results. The absence of prices in the data forces me to conflate "revenue productivities" with "quantity productivity" (Foster, Haltiwanger, and Syverson 2008; Hsieh and Klenow 2009). Dollar values of productivity—"revenue productivities"—can change as a result of movements in prices as well as in physical productivities. The former, in turn, can be driven by shifts in entry barriers, independently of any changes in physical productivity.

Consider, for example, an economy in which the failure of marginal value labor productivities to equalize across sectors is due to barriers to labor mobility. As in Hsieh and Klenow (2009), I could suppose these barriers are government-imposed restrictions, such as preferential licensing or credit policies that advantage some firms or sectors over others. In this world, differences in dollar labor productivities within countries would reflect the magnitude of these barriers, and my regressions would be capturing convergence and divergence in them.[16] In particular, the "unconditional convergence" result may reflect the fact that industries with high barriers have experienced a decline in barriers while industries with low barriers have experienced an increase. Interesting as this reading of the evidence may be in its own right, it would be quite a different result than one about productivity convergence.

However, in this case I am looking at industries across as well as within countries. In fact, as the industry-by-industry regressions discussed previously (and the sigma-convergence results, discussed later) indicate, an essential part of the identification for unconditional convergence comes from the variation within industries across different countries. Revenue productivities are converging *globally*. It is not very plausible to attribute differences in dollar labor productivities in, say, motor vehicles across countries to differences in entry barriers faced by motor vehicle producers in different countries, or to attribute convergence in revenue productivities to systematic changes in such barriers.

In addition, because low-productivity industries are in poor countries, for changes in labor mobility barriers to account for my findings, mobility barriers must have come down in poorer countries and increased in the rich countries. If this were true, I would also see much more rapid expansion in manufacturing in poor countries. As I show later when I discuss aggregation issues, this doesn't seem to be the case. The movement of labor into the manufacturing industries of the poorer countries is not significant or rapid enough to support such a conjecture. The sluggish pace of labor movement, in particular the relatively slow expansion of manufacturing in the poorest countries, does not suggest systematic changes in barriers to mobility.

16. I am grateful to a referee for raising this possibility and for the related discussion on revenue versus quantity productivity.

III.D. Sigma-Convergence

Beta-convergence does not guarantee sigma-convergence in a world where growth rates are driven not just by the forces of convergence but also by other determinants and shocks. The present data do not allow a very comprehensive analysis of sigma-convergence because I need to have a large sample of countries with data both at the beginning and end of the period to determine whether the dispersion of productivity has diminished over any given time horizon. As discussed previously, the country coverage shrinks dramatically when I restrict the sample to a cross-section of any fixed time period.

The best I can do is to choose a time period that maximizes the number of countries that can be included. Table V shows the results for the 1995–2005 decade, which covers a variable number of countries across different industries, and 63 countries for manufacturing as a whole. The overall pattern across two-digit manufacturing industries seems mixed, but a majority (16) shows declines in dispersion, some quite dramatically. Basic metals (ISIC 27); office, accounting, and computing machinery (ISIC 30); and radio, television, and communication equipment (ISIC 32) have all experienced substantial reductions in the standard deviation of log productivity, of the order of 30 log-points or more, in a period as short as a decade. More remarkably, dispersion in manufacturing as a whole has come down by 10 log-points in my sample of 63 countries. This amounts to a reduction of 10% in dispersion.

Figures VI and VII provide a visual sense of these findings. The distribution of aggregate manufacturing productivity has a clear twin-peaked shape. Between 1995 and 2005, the two peaks have moved considerably closer to each other (Figure VI).

IV. Why Unconditional Convergence Does Not Aggregate Up

The forces of convergence seem quite strong in manufacturing industries. It stands to reason that one ould uncover similar results for certain other parts of the economy as well, perhaps modern, tradable services such as financial or business services, among others. One might expect convergence at the sectoral level to produce aggregate convergence as well, unless there are countervailing forces pushing in the other direction. Yet the aggregate

TABLE V

SIGMA-CONVERGENCE, 1995–2005

ISIC code	Sector	Number of countries	Sigma in 1995	Sigma in 2005	Difference
D	Total manufacturing	63	1.186	1.082	−0.104
15	Food and beverages	43	1.086	0.942	−0.144
16	Tobacco products	31	1.591	1.668	0.077
17	Textiles	54	1.218	1.246	0.028
18	Wearing apparel, fur	34	1.166	1.301	0.135
19	Leather, leather products, and footwear	34	1.144	1.269	0.125
20	Wood products (excl. furniture)	53	1.479	1.154	−0.325
21	Paper and paper products	52	1.188	0.999	−0.189
22	Printing and publishing	53	1.167	1.060	−0.107
23	Coke, refined petroleum products, nuclear fuel	38	1.415	1.835	0.420
24	Chemicals and chemical products	52	1.286	1.273	−0.013
25	Rubber and plastics products	53	1.099	0.976	−0.123
26	Nonmetallic mineral products	56	1.206	0.992	−0.214
27	Basic metals	50	1.291	0.893	−0.398
28	Fabricated metal products	53	1.231	1.108	−0.123
29	Machinery and equipment n.e.c.	37	1.237	1.247	0.010
30	Office, accounting, and computing machinery	23	0.994	0.688	−0.306
31	Electrical machinery and apparatus	32	1.124	0.943	−0.181
32	Radio, television, and communication equipment	26	1.229	0.914	−0.315
33	Medical, precision, and optical instruments	45	1.380	1.202	−0.178
34	Motor vehicles, trailers, semi-trailers	35	0.945	1.012	0.067
35	Other transport equipment	27	1.092	0.944	−0.148
36	Furniture; manufacturing n.e.c.	52	1.257	1.135	−0.122
37	Recycling	22	0.991	0.988	−0.003

FIGURE VI

Sigma-Convergence, All Manufacturing

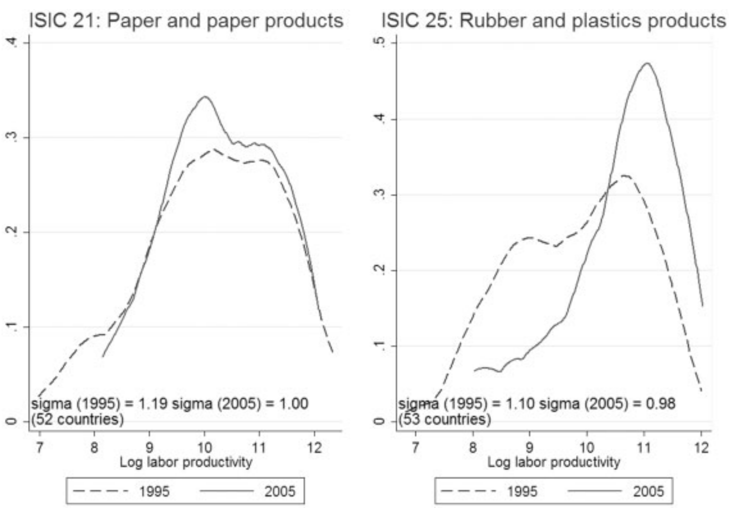

FIGURE VII

Sigma-Convergence in Specific Industries data do not support this conjecture. In this section I consider why economies as a whole fail to exhibit unconditional converge despite the strong pull of convergence within manufacturing industries.

Recall (from Table II) that aggregate manufacturing does exhibit unconditional convergence. There is some evidence that convergence gets stronger the more one disaggregates within manufacturing. But it is clear that the bulk of the convergence failure takes place as one goes from manufacturing (in aggregate) to the rest of the economy. So I focus here on the contrasting behaviors of manufacturing and the economy (both as a whole), abstracting from aggregation issues within manufacturing.

To compare manufacturing to the rest of the economy, I need data that go beyond what I have been using so far. I proceed by combining INDSTAT2 with Penn World Tables 7.0 (Heston, Summers, and Aten 2011; PWT), which include data on aggregate GDP and (implicitly) total employment. Because data for GDP are in real purchasing power parity (PPP)-adjusted terms in the PWT, I first convert these to current dollars to render the economywide data directly comparable to INDSTAT2.[17] Subtracting INDSTAT2 manufacturing value added and employment levels from the aggregates in PWT yields presumptive values for nonmanufacturing. We thereby obtain labor productivity figures for nonmanufacturing and the entire economy that are consistent with both data sets. The employment shares of manufacturing and nonmanufacturing can be imputed in a similar fashion.[18]

17. PWTs include data on real GDP per worker (*rgdpwok*), which is PPP-converted GDP chain per worker at 2005 constant prices. To undertake the conversion, first I recover the conversion factor between current and constant prices, by taking the ratio of current and constant price GDP data in PWT. For current GDP data, I use PPP-converted GDP per capita, G-K method, at current prices (in I$) (*cgdp*). For constant price GDP data, I use PPP-converted GDP per capita (Laspeyres), derived from growth rates of c, g, i, at 2005 constant prices (*rgdpl*). The ratio of these two values ($\frac{cgdp}{rgdpl}$) gives a price conversion factor between current and constant. Next I calculate a conversion factor between market prices and PPP. For this I use the PWT values for PPP over GDP in national currency units per US$ (*ppp*) divided by the exchange rate to US$ in national currency units per US$ (*xrat*). Using these two conversion factors, I convert PWT's data on real GDP per worker into nominal U.S. dollars ($ngdpwok = rgdpwok * (\frac{cgdp}{rgdpl}) * (\frac{pp}{xrat})$). This gives nominal GDP chain per worker at current prices (US$).

18. To compute the employment share of manufacturing, I first compute the total working population using data from PWTs. To do this, I divide PWT data on

I can now compare convergence behavior systematically across different types of activities. I focus in this section on the 1995–2005 cross-section, because these regressions do not mix different time periods and are the easiest ones to interpret in the present context. Column (2) of Table VI verifies the absence of unconditional convergence in nonmanufacturing. The estimated convergence coefficient is very small and statistically indistinguishable from zero. Column (3) replicates the exercise for economywide labor productivity, again confirming nonconvergence. For comparison purposes, the bottom part of the table shows the analogous results for the full panel.[19]

Other columns in Table VI highlight the importance of the relative size of manufacturing in driving convergence behavior. The employment share of manufacturing (α) appears to be a key conditioning factor. As columns (5)–(7) show, α is a significant determinant of economywide growth. More important for my purposes, aggregate convergence seems to be conditional on α. Comparing columns (3) and (7), one can see that initial economywide productivity turns statistically significant in the aggregate growth equation once I control for α. These are, of course, sparse

GDP per capita (*rgdpch*) by data on GDP per worker (*rgdpwok*). This gives the number of workers per capita. From this number and the total population figures (*pop*) in PWT, I calculate total employment. Total manufacturing employment is given by INDSTAT2 as the number of workers in manufacturing as a whole, in ISIC category D. From this and the total employment number computed using PWT, the employment share of manufacturing (α) can be calculated. To compute nonmanufacturing labor productivity, I convert PWT data on total PPP converted GDP, G-K method, at current prices in millions I$ (*tcgdp*) to nominal GDP, using the PPP–exchange rate conversion factor described in the previous note ($\frac{ppp}{xrat}$). Nonmanufacturing value added is calculated as the difference between this nominal value and the value added for manufacturing as a whole from INDSTAT2, in ISIC category D. This number is then divided by nonmanufacturing employment (the difference between total employment and total manufacturing employment), to give nonmanufacturing labor productivity. From these numbers, growth rates can be calculated to run convergence regressions for nonmanufacturing.

19. Interestingly, nonmanufacturing and the aggregate economy appear to exhibit some unconditional convergence in the baseline sample, at a rate of around 1%, or one third that for manufacturing in the same sample. However, this is not robust and seems to be a result of the correlation between the initial levels of productivity in manufacturing and nonmanufacturing. When initial manufacturing productivity is also included in the regression, the coefficient on initial nonmanufacturing productivity actually turns positive in both the aggregate economy and nonmanufacturing growth regressions (while manufacturing productivity is negative and significant).

TABLE VI

ECONOMIC STRUCTURE AND AGGREGATE (NON)COVERGENCE

	Dependent variable: decadal growth of labor productivity in						
	(1) Manufacturing	(2) Nonmanufacturing	(3) Economy	(4) Economy	(5) Economy	(6) Economy	(7) Economy
Panel A: 1995–2005 cross-section							
Log initial productivity, manufacturing	−0.020***			−0.026***	−0.015*	−0.013***	
	(0.006)			(0.009)	(0.008)	(0.004)	
Log initial productivity, nonmanufacturing		−0.002		0.015**	0.002		
		(0.003)		(0.006)	(0.007)		
Log initial productivity, economy (GDP/worker)			−0.003				−0.010***
			(0.003)				(0.003)
Alpha (employment share of manufacturing)					0.290***	0.311***	0.386***
					(0.093)	(0.071)	(0.093)
Period fixed effects	yes	yes	yes	yes	yes	yes	yes
Number of countries	61	61	61	61	61	61	61
Number of observations	61	61	61	61	61	61	61
Panel B: Panel specification							
Log initial productivity, manufacturing	−0.010***			−0.023***	−0.010*	−0.011***	
	(0.003)			(0.005)	(0.005)	(0.003)	
Log initial productivity, nonmanufacturing		0.002		0.016***	-0.002		
		(0.002)		(0.004)	(0.005)		
Log initial productivity, economy (GDP/worker)			0.003				−0.009***
			(0.002)				(0.003)
Alpha (employment share of manufacturing)					0.309***	0.297***	0.344***
					(0.049)	(0.034)	(0.043)
Period fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of countries	95	95	95	95	95	95	95
Number of observations	234	234	234	234	234	234	234

Notes. Country sample is limited to those in the 1995–2005 cross-section and panel samples with data for all columns. Panel specification pools data for 1965–75, 1975–85, 1985–95, and 1995–2005. Robust standard errors (Panel A) and standard errors clustered at country level (Panel B) are reported.

***$p < .01$, **$p < .05$, *$p < .10$.

specifications that do not rule out alternative interpretations; α could be proxying for a wide range of other conditioning variables with which it is correlated. The "conditional" convergence rate controlling for α is relatively low (1%) in light of estimates in Barro (2012), so this proxying is likely imperfect. In any case, these results are suggestive of the role played by the relative size of manufacturing in explaining convergence behavior across countries.

To investigate the issue more formally, I divide the economy into manufacturing (m) and nonmanufacturing (n) activities. GDP per worker is the weighted average of labor productivity in these two activities: $y = \alpha y_m + (1-\alpha)y_n$, where the weight α is the share of the economy's labor force employed in manufacturing. (I dropped country subscripts to avoid clutter.) Growth in GDP per worker is in turn expressed as

$$\hat{y} = \alpha\theta_m \hat{y}_m + (1-\alpha)\theta_n \hat{y}_n + (\theta_m - \theta_n)d\alpha,$$

where a "^" over a variable denotes proportional growth rates as before, and $\theta_m = \frac{y_m}{y}$ and $\theta_n = \frac{y_n}{y}$ are the productivity premia/discounts for the two sectors.

I now impose some further structure on the growth decomposition, using the convergence results obtained thus far. In particular, I write the growth rates of manufacturing and nonmanufacturing as:

$$\hat{y}_n = g$$

$$\hat{y}_m = g + \beta\left(\ln y^* - \ln y_m\right),$$

where g is the underlying long-term balanced growth rate of the economy, y^* is the productivity frontier in manufacturing, y_m is manufacturing labor productivity in the home economy, and β (>0) is the convergence coefficient in manufacturing. This formulation captures the basic asymmetry between manufacturing and nonmanufacturing, namely, that manufacturing is the beneficiary of a convergence "kick," which peters out as the economy gets closer to the frontier.

Substituting these in and noting that $\alpha\theta_m + (1-\alpha)\theta_n = 1$, economywide growth becomes

$$\hat{y} = g + \alpha\theta_m\beta(\ln y^* - \ln y_m) + (\theta_m - \theta_n)\, d\alpha. \tag{4}$$

So growth equals an exogenous (or country-specific) component, a manufacturing convergence factor (that is decreasing in the level of manufacturing productivity), and a reallocation term. The reallocation term captures the effect of changes in the composition of employment across sectors when productivities differ between manufacturing and nonmanufacturing. In particular, because $\theta_m > \theta_n$ in the data, an increase in manufacturing employment share ($d\alpha$) raises growth. I denote the reallocation term as $\Delta \equiv (\theta_m - \theta_n)\, d\alpha$.

Equation (4) is helpful for understanding why manufacturing convergence does not translate into aggregate convergence. Most critically, it highlights the role of α, the share of manufacturing employment. The economywide impact of manufacturing growth is mediated through this variable. One characteristic of poor countries is that they have very small formal manufacturing sectors. The average α for the poorest half of the baseline sample is only around 5%. This means that even powerful convergence effects for manufacturing have only tiny consequences for the economy as a whole. The manufacturing premium θ_m tends to be larger in poorer countries, but in practice this only partly compensates for the lower α (see later discussion).

In light of this, I would need the reallocation term Δ to not only be large but to also vary systematically with incomes, with poorer countries benefiting substantially more from reallocation towards manufacturing. However, I fail to find such strong and systematic reallocation effects in the data.

I illustrate these arguments with a numerical exercise, which quantifies equation (4). Table VII splits the 1995–2005 sample into 10 deciles according to initial level of aggregate GDP per worker. I then compute the terms in equation (4) for each decile. To perform the calculations, I choose values for α, θ_m, and y_m that correspond to the averages for each decile. For β, I use .020, which is the value estimated for this sample of observations.[20] For y^* I use the average manufacturing productivity in the top decile of the sample (for manufacturing productivity). The value of g is set equal to 0 with no loss of generality. For each decile, the table shows the predicted manufacturing growth rate $\beta(\ln y^* - \ln y_m)$, the predicted aggregate growth rate due to manufacturing convergence $\alpha\theta_m\beta(\ln y^* - \ln y_m)$, and

20. This is different from the estimate noted in Table II because I am working with the 1995–2005 sample here (and not the baseline sample).

the total predicted growth rate including the reallocation term $\alpha\theta_m\beta(\ln y^* - \ln y_m) + \Delta$.[21]

For example, for the poorest 10% of the sample, the average values of the parameters are as follows: $\alpha = 0.0278$, $\theta_m = 7.7451$, and $\ln y^* - \ln y_m = 3.5160$. As the first row of Table VII shows, this yields a predicted manufacturing growth rate of 7.1%, an aggregate convergence term of 0.7%, and a predicted overall growth rate that is actually slightly lower at 0.6%.

Table VII shows a steep negative gradient in manufacturing growth as initial GDP per worker increases (column (4)). This is as expected from the results presented previously. The predicted manufacturing growth rate goes down from 7.1% for the bottom decile to 0.6% at the very top.[22] The table also shows that very little of this gradient survives at the aggregate level once the effect is scaled down by $\alpha\theta_m$, which is substantially smaller than 1. The aggregate convergence term is a mere 0.7% for the bottom decile (an order of magnitude smaller). The key contributor here is α, which not only is very small but increases as incomes rise, further moderating the forces of convergence (column (1)). This is only partly offset by the fall in θ_m (column (2)).

Furthermore, the reallocation term is tiny, often negative, and does not vary substantially across income levels to make much difference (column (6)). In principle, this term could have made a substantial contribution to convergence if labor were to move more rapidly to manufacturing in response to the large return differentials across sectors. The labor productivity differential between manufacturing and nonmanufacturing ($\theta_m - \theta_n$) is of the order of 300–700% for the lowest deciles. With a differential of 500%, even if α were to increase by 0.5% annually (resulting in roughly a 5 percentage point increase in the employment share of manufacturing over a decade), the poorest economies in the sample would experience a growth boost of around $5 \times 0.5 = 2.5$ percentage points. As it is, the actual change in α is a minute fraction of that and often negative (as can be observed from the

21. I use end-of-period θ_m and θ_n to compute Δ because I am dealing with discrete changes.

22. There is a small convergence effect even in the top decile because manufacturing convergence depends on distance from the *manufacturing* frontier. Table VI organizes deciles according to GDP per worker, not manufacturing productivity.

TABLE VII

PREDICTED GROWTH RATES BY INCOME (GDP/WORKER) DECILE

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	α	θ_m	$\ln y^* - \ln y_m$	Manufacturing growth $\beta\times(3)$	Aggregate convergence term $(1)\times(2)\times(4)$	Δ	Predicted aggregate growth $(5)+(6)$
Bottom decile	0.0278	7.7451	3.5160	7.09%	0.67%	−0.0016	0.51%
Decile 9	0.0739	3.4913	3.3287	6.71%	0.84%	−0.0053	0.32%
Decile 8	0.0488	4.3041	3.0440	6.14%	0.42%	−0.0004	0.38%
Decile 7	0.0731	3.7142	2.5166	5.07%	1.06%	−0.0005	1.01%
Decile 6	0.0964	2.5252	2.3940	4.83%	0.69%	−0.0003	0.66%
Decile 5	0.0794	2.8562	1.8194	3.67%	0.62%	−0.0007	0.55%
Decile 4	0.1307	2.0218	1.7922	3.61%	0.62%	−0.0003	0.60%
Decile 3	0.1415	1.3768	1.1459	2.31%	0.37%	−0.0007	0.31%
Decile 2	0.1385	1.4440	0.7061	1.42%	0.17%	0.0001	0.18%
Top decile	0.1541	1.2064	0.2909	0.59%	0.09%	−0.0006	0.03%

Notes. Sample: 1995–2005, $\beta = .020$, $\ln y^* = 11.5098$; α, θ_m, $\ln y^*$ and $\ln y_m$ are beginning-of-period values. Δ is calculated by multiplying the change in α over the decade with the end-of-period differences in θ's, and dividing the result by 10 to annualize it.

negative entries in column (6)).[23] Very similar results are obtained for the panel and baseline samples as well (not shown).

The net result is that the gradient for predicted economywide growth is barely distinguishable from zero (see column (7)). The difference between the gradients for manufactures and the aggregate economy is illustrated in Figure VI. In fact, the implied convergence factor for economywide growth is so small is that it would be easily swamped by random measurement error or unobserved country-level determinants. For example, the small negative gradient in Table VII for predicted overall growth becomes statistically indistinguishable from zero if I augment equation (4) with a random error term distributed normally with mean 0 and variance 0.0001 (a standard deviation in unexplained growth rates across countries of just 1 percentage point).

In sum, aggregate nonconvergence appears to be explained by the following combination of facts: (1) non-manufacturing activities do not exhibit unconditional convergence; (2) poor countries have little employment in manufacturing, depressing the contribution of manufacturing to overall growth; (3) the share of employment in manufacturing rises over the course of development, giving less-poor countries a growth boost; and (4) the reallocation effect is neither sizable enough nor systematically larger at lower income levels. In terms of quantitative magnitudes, the first two factors play the dominant roles. From an economic standpoint, however, the last fact is perhaps most interesting, pointing to an important unexploited potential in poor countries.

V. Concluding Remarks

I have provided evidence in this article that unconditional convergence is alive and well. One needs to look for it among manufacturing industries rather than entire economies. It is perhaps not surprising that manufacturing industries should exhibit unconditional convergence and, if the estimates here are to be believed, at quite a rapid pace, too. These industries produce tradable goods and can be rapidly integrated into global production

23. Wong (2006) finds a very small reallocation effect in his study of 13 OECD countries as well. Wong performs an accounting decomposition for OECD economies, allocating convergence among seven different sectors and an interaction (reallocation) term.

FIGURE VIII

Difference between Manufacturing and Economy-wide Growth Gradients. See text for explanation.

networks, facilitating technology transfer and absorption. Even when they produce just for the home market, they operate under competitive threat from efficient suppliers from abroad, requiring that they upgrade their operations and remain efficient. Traditional agriculture, many nontradable services, and especially informal economic activities do not share these characteristics.

The findings in this article offer new insight on the determinants of economic growth and convergence across countries. They suggest that lack of convergence is due not so much to economy-wide misgovernance or endogenous technological change but to specific circumstances that influence the speed of structural reallocation from nonconvergence to convergence activities. The policies that matter are those that bear directly on this reallocation. As discussed in McMillan and Rodrik (2011) and Rodrik (2011b), what high-growth countries typically have in common is their ability to deploy policies that compensate for the market and government failures that block growth-enhancing structural transformation. Countries that manage to affect the requisite structural change grow rapidly, and those that fail don't.

Put differently, successful countries experience both productivity convergence in formal manufacturing and rapid industrialization. Unsuccessful countries make do with just the former.

Supplementary Material

An Online Appendix for this article can be found at QJE online (qje.oxfordjournals.org).

John F. Kennedy School of Government, Harvard University

REFERENCES

Acemoglu, Daron. *Introduction to Modern Economic Growth* (Princeton, NJ: Princeton University Press, 2009).

Barro, Robert J., "Convergence and Modernization Revisited," NBER Working Paper No. 18295, 2012.

Barro, Robert J., and Xavier Sala-i-Martin, "Convergence across States and Regions," *Brookings Papers on Economic Activity*, 1 (1991), 107–158.

Baumol, William, "Productivity Growth, Convergence, and Welfare: What the Long-Run Data Show," *American Economic Review*, 76, no. 5 (1986), 1072–1085.

Bénétrix, Agustín S., Kevin H. O'Rourke, and Jeffrey G. Williamson, "The Spread of Manufacturing to the Periphery 1870–2007: Eight Stylized Facts," NBER Working Paper No. 18221, 2012.

Bernard, Andrew, and Charles Jones, "Productivity and Convergence across U.S. States and Industries," *Empirical Economics*, 21 (1996a), 113–135.

———, "Comparing Apples to Oranges: Productivity Convergence and Measurement across Industries and Countries," *American Economic Review*, 86 (1996b), 1216–1238.

DeLong, J. Bradford, "Productivity Growth, Convergence, and Welfare: Comment," *American Economic Review*, 78, no. 5 (1988), 1138–1154.

Durlauf, Steven, Paul Johnson, and Jonathan Temple, "Growth Econometrics," in *Handbook of Economic Growth*, P. Aghion and S. Durlauf, eds. (Amsterdam: North-Holland, 2005).

Foster, Lucia, John Haltiwanger, and Chad Syverson, "Reallocation Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98, no. 1 (2008), 394–425.

Heston, Alan, Robert Summers, and Bettina Aten, *Penn World Table Version 7.0*. Center for International Comparisons of Production, Income and Prices, University of Pennsylvania, May 2011.

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

Hurwicz, Leonid, "Least-Squares Bias in Time Series," in *Statistical Inference in Dynamic Economic Models*, Tjalling C. Koopmans, ed. (New York: Wiley, 1950).

Hwang, Jason J., "Patterns of Specialization and Economic Growth," Ph.D dissertation, Harvard University, 2007.

Levchenko, Andrei A., and Jing Zhang, "The Evolution of Comparative Advantage: Measurement and Welfare Implications," NBER Working Paper No. 16806, 2011.
McMillan, Margaret, and Dani Rodrik, "Globalization, Structural Change, and Productivity Growth," NBER Working Paper No. 17143, 2011.
Nickell, Stephen, "Biases in Dynamic Models with Fixed Effects," *Econometrica*, 49 (1981), 1417–1426.
Rodrik, Dani, "Unconditional Convergence," NBER Working Paper No. 17546, 2011a.
———, "The Future of Economic Convergence," NBER Working Paper No. 17400, 2011b.
Schott, Peter K., "Across-Product versus Within-Product Specialization in International Trade," *Quarterly Journal of Economics*, 119 (2004), 647–678.
Sørensen, Anders, "Comparing Apples and Oranges: Productivity Convergence and Measurement across Industries and Countries: Comment," *American Economic Review*, 91, no. 4 (2001), 1160–1167.
Subramanian, Arvind. *Eclipse: Living in the Shadow of China's Economic Dominance* (Washington, DC: Peterson Institute for International Economics, 2011).
Temple, Jonathan R. W., "Robustness Tests of the Augmented Solow Model," *Journal of Applied Econometrics*, 13, no. 4 (1998), 361–375.
United Nations Industrial Development Organization (UNIDO), *INDSTAT4 Industrial Statistics Database*, 2011.
———, *INDSTAT2 Industrial Statistics Database*, 2012.
Wong, Wei-Kang, "OECD Convergence: A Sectoral Decomposition Exercise," *Economics Letters*, 93 (2006), 210–214.
Young, Andrew T., Matthew J. Higgins, and Daniel Levy, "Sigma Convergence versus Beta Convergence: Evidence from U.S. County-Level Data," *Journal of Money, Credit and Banking*, 40, no. 5 (2008), 1083–1093.