

The Productivity of Nations

Margarida Duarte and Diego Restuccia

In this article, we document observations on labor productivity across countries over time. Using data from Heston, Summers, and Aten (2002) on gross domestic product (GDP) per worker—our measure of labor productivity—we emphasize three main facts about the distribution of labor productivity across countries between 1960 and 1996.¹ First, there is substantial dispersion in labor productivity across countries. For instance, in 1960 an average worker in the richest 5 percent of countries in the world produces about 35 times more output than an average worker in the poorest 5 percent of countries in the world.

Second, disparity in labor productivity has increased over time. By 1996, the labor productivity ratio between the richest and poorest countries increased to approximately 46. This increase in disparity is explained by a substantial deterioration in labor productivity in the poorest countries of the world relative to that of the United States. We report several statistics of dispersion indicating that labor productivity differences between the richest and poorest countries have increased since the mid-1980s. This characterization of increased dispersion in labor productivity contrasts with a relative stability documented in previous studies, in which the coverage period ended in 1985.

Third, there is substantial mobility of individual countries in the distribution of labor productivity over time. For instance, labor productivity in Hong Kong relative to that in the United States rose from 19 percent in 1960 to 94 percent in 1996—an increase of a factor of almost 5 during the period—while relative labor productivity in Venezuela declined from 94 percent in 1964 to 36 percent in 1996—a more than twofold drop in relative productivity. We also document a number of individual episodes of growth and decline. Ac-

■ We would like to thank Borys Grochulski, Andreas Hornstein, and Leonardo Martinez for their comments and Andrea Waddle for her comments and excellent research assistance. All errors are our own. The views expressed in this article are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

¹ Throughout the article, we refer to GDP per worker, output per worker, and labor productivity interchangeably.

counting for these specific labor productivity paths may prove useful not only from a policy perspective, but also in testing and improving existing theories of productivity levels or in the development of alternative theories.

This article relates to a large literature on the world distribution of income that includes Kaldor (1961), Kuznets (1966), Maddison (1995, 2001), Parente and Prescott (1993), Chari, Kehoe, and McGrattan (1996), Jones (1997), among many others. We contribute to this literature by adding the period between 1985 and 1996 to the analysis. In addition, we focus on output per worker as opposed to output per capita since output per worker relates more directly to theories of labor productivity. The difference between the two measures is given by the employment-to-population ratio. While there are substantial differences in these ratios across countries, the differences are not systematically related to development. Therefore, our summary statistics characterizing output per worker over time are similar to statistics calculated using output per capita. However, there are substantial changes in employment-to-population ratios for individual countries over time, and these changes are not systematically related to development or growth in relative productivity. Therefore, for an individual country, changes in output per capita can severely overstate or understate changes in labor productivity.

There are two additional differences between this article and the previous literature. First, we characterize disparity and mobility using trended data. That is, we use the Hodrick-Prescott filter to abstract from business-cycle fluctuations in the data. Second, we seek to systematically identify remarkable episodes of growth (positive or negative) in the data at some point during the 1960–1996 period. In the literature, countries facing these episodes are typically referred to as miracles and disasters. We document 13 miracle and 17 disaster episodes in our data set. Among the miracle episodes, we report the movement of labor productivity in Botswana relative to that in the United States from 7 percent to 30 percent in 26 years; in Hong Kong, from 19 percent to 94 percent in 36 years; and in China, from 4 percent to 8 percent in 18 years. We also document the recent, but not yet as long, growth episodes of Chile, Ireland, and India, which may become miracle episodes within the next two decades.

Furthermore, we also systematically document depression episodes in our panel data. Ever since the study of the Great Depression in the United States by Cole and Ohanian (1999), there has been substantial interest in studying depression episodes (defined broadly as periods of lower-than-usual relative productivity).² We follow this literature in characterizing depression episodes by using the raw data on output per worker relative to a trend growth of

² This marked interest is reflected, for instance, in the work of Prescott (2002) and in several articles published in a special volume of the *Review of Economic Dynamics* edited by T. Kehoe and E.C. Prescott (see Kehoe and Prescott 2002).

2 percent per year. Even in our relatively small sample period, we find that depressions are quite common, both among rich and poor countries. We report 29 depression episodes in Section 4.

Our study also relates to a broad literature on models that seek to explain development facts such as those of disparity and mobility discussed above. For excellent surveys of this literature see, for instance, Klenow and Rodríguez-Clare (1997) and Caselli (2005). Our study complements this literature by expanding and updating the set of facts that theories of development should be able to explain.

This article is organized as follows. In the next section, we describe in detail the data we use for the analysis. Section 2 documents the main facts about dispersion and mobility in the distribution of labor productivity. In Section 3, we discuss the remarkable episodes of growth during the period, namely miracle and disaster episodes, and in Section 4, we document the episodes of depressions. Section 5 discusses our main findings relative to those using instead output per hour, an alternative sample of countries, and output per capita. We conclude in Section 6.

1. DATA

We focus on output per worker as our measure of labor productivity. We use annual data on PPP-adjusted GDP per worker in chained 1996 prices obtained from Heston et al. (2002), also known as the Penn World Table V6.1 (PWT6.1). We choose the PWT6.1 for our analysis because it is the most comprehensive source of comparable measures of output per worker across countries.

We focus on output per worker in order to emphasize the connection of the data with research on productivity differences across countries. Although output per hour is a more complete measure of labor productivity, we abstract from differences in hours per worker across countries due to the lack of systematic data for a large number of countries over the entire period of our study. However, in Section 5, we use the available data on hours per worker to calculate output per hour and discuss our findings relative to this more complete measure of labor productivity. Measures of output per capita are appropriate when the focus is on wealth differences across countries.³ We also discuss our findings relative to the use of output per capita in Section 5.

³ Maddison (1995, 2001) documents comparable measures of output per capita for a wide range of countries and time periods. However, we note that, as in the PWT6.1, Maddison uses the detailed price data from the International Comparisons Project (ICP) of the United Nations to calculate purchasing power parity (PPP) conversion factors at a point in time, and national accounts data to extrapolate over time. In this sense, Maddison's output data is comparable to PWT6.1, especially for the Benchmark countries (see Section 5 for a definition of Benchmark countries).

Our data set consists of annual observations for 99 countries from 1960 to 1996.⁴ The countries in our data set satisfy two restrictions. First, the total population of each country was at or above 1 million people in 1996. Second, data was available at each date from 1960 to 1996. We make this second restriction in order for the sample of countries to be constant throughout the period of analysis. Data is available for 48 countries from 1950 to 2000. However, countries without observations between 1950 and 1960 and between 1996 and 2000 tend to be poorer countries. Thus, the smaller sample from 1950 to 2000 is less representative of the world distribution of output per worker than our sample with 99 countries from 1960 to 1996.

In documenting observations about dispersion, mobility, and miracles and disasters, we abstract from business-cycle fluctuations and trend the data using the Hodrick-Prescott filter.⁵ Abstracting from business-cycle fluctuations when reporting development facts is not innocuous. As we document in Section 4, countries undergo episodes of substantial growth and decline that are not entirely related to their development process. To illustrate the cycles in the annual data, we report in Figure 1 the raw data on output per worker and the trended data for four countries: the United States, Argentina, Romania, and Switzerland. As is true with many other countries in our panel data, these countries have undergone relatively short-lived variations in their output per worker at different points in time. Our documentation of the development facts abstracts from these fluctuations.

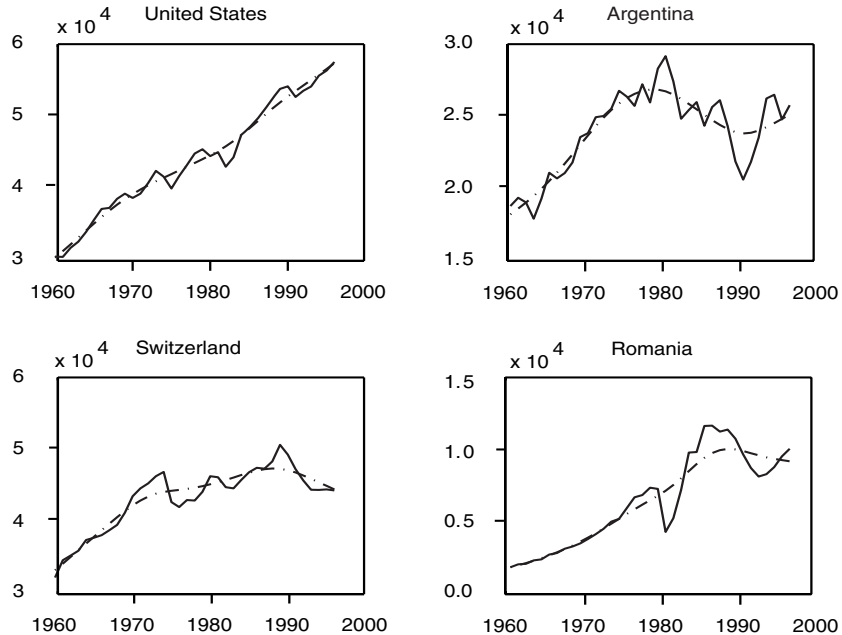
For the most part, we report statistics on output per worker relative to that of the United States. Our view is that the United States is a rich, stable, and diverse country. For most of the period of analysis, the United States had the highest labor productivity. Moreover, in the post-war period, (trended) labor productivity grew at roughly 2 percent per year. Therefore, the United States represents a good benchmark against which to measure potential gains in labor productivity in all countries.

2. LABOR PRODUCTIVITY ACROSS COUNTRIES

We emphasize three facts about the distribution of labor productivity across countries. First, there is a large disparity in output per worker across countries. Second, there is a substantial increase in disparity over time. Third, there are substantial movements of individual countries in the world distribution of productivity. In the remainder of this section we characterize these facts in detail.

⁴ See the Appendix for a list of countries in our data set.

⁵ We set the smoothing parameter λ equal to 100.

Figure 1 Output per Worker in Four Countries

Notes: For each country, the solid line represents the raw data on output per worker and the dotted-dashed line the HP-trended data.

Disparity


A remarkable fact of modern development data is the large disparity in productivity among countries. Here we focus on different measures of disparity and their evolution between 1960 and 1996.

We start by focusing on the five richest and five poorest countries in our sample. We compute the ratio of average output per worker for the five richest and five poorest countries for each year from 1960 to 1996, illustrated in Figure 2. This ratio varies between 35 and 46 over the period of analysis. That is, the average worker in the richest countries produces between 35 and 46 times more output than the average worker in the poorest countries. These are remarkable differences in labor productivity.

This measure of disparity in productivity across countries has been roughly constant from 1960, at 35, until the mid-1980s. The ratio declined slightly around 1980 but has increased steadily since then to a factor of 46 in 1996.⁶

⁶ The ratio of average output per worker for the 5 richest and 5 poorest countries computed using the data set of 48 countries with data from 1950 to 2000 shows the same pattern as in

Figure 2 Output per Worker—Ratio of Five Richest to Five Poorest Countries



Notes: Between 1960 and 1996, the following countries comprised the five richest at some point in time: the United States, Switzerland, Canada, New Zealand, the Netherlands, Belgium, Italy, Norway, and Hong Kong. During the same period, the following countries comprised the five poorest countries at some point in time: Tanzania, Guinea Bissau, Burundi, Ethiopia, Burkina Faso, Uganda, and the Democratic Republic of Congo.

This increase in dispersion runs contrary to the established view in the development literature that dispersion in the world distribution of productivity has been rather constant over time (see, for instance, Parente and Prescott 1993; Chari, Kehoe, and McGrattan 1996). The reason for this view is that until about 1985 (the end date in most previous studies), the productivity ratio of the richest to the poorest countries was roughly constant. See Figure 2 where the line drawn at 1985 emphasizes the connection with the earlier literature.

In Figure 3 we report the relative productivity of the five poorest and five richest countries between 1960 and 1996, each normalized to 100 in

Figure 2. For this data set, the ratio of rich to poor increases steadily from the early 1980s to 2000.

