
Overview

By the problem of economic development I mean simply the problem of accounting for the observed pattern, across countries and across time, in levels and rates of growth of per capita income. This may seem too narrow a definition, and perhaps it is, but thinking about income patterns will necessarily involve us in thinking about many other aspects of societies too, so I would suggest that we withhold judgement on the scope of this definition until we have a clearer idea of where it leads us.

Robert E. Lucas (1988)

[W]e should never lose sight of the ultimate purpose of the exercise, to treat men and women as ends, to improve the human condition, to enlarge people's choices. [A] unity of interests would exist if there were rigid links between economic production (as measured by income per head) and human development (reflected by human indicators such as life expectancy or literacy, or achievements such as self-respect, not easily measured). But these two sets of indicators are not very closely related.

Paul P. Streeten (1994)

2.1. Introduction

Economic development is *the* primary objective of the majority of the world's nations. This declaration is uncontroversial, or so it would appear in public discourse at least. Every year, elaborate plans are hatched, investments are undertaken, policies are devised, and government budgets are framed to achieve this goal, or at least to get closer to it. How do we identify and track the results of these efforts? What criteria do we use to evaluate the extent of "development" a country has undergone or how "developed" or "underdeveloped" a country is at present? How do we measure development?

These are not easy questions. We all have intuitive notions of development. Presumably, when we speak of a developed society, we have in mind a world in which people are well fed and well clothed, have access to a variety of goods and services,

possess the luxury of leisure and entertainment, and live in a healthy environment. We think of a society free of discrimination or violence, with reasonable economic equality, where the sick receive proper medical care and people do not have to sleep on the sidewalks. In short, most of us would insist that a *minimal* requirement for a developed nation is that its *physical* quality of life be high, and be uniformly so rather than restricted to an incongruously affluent minority.

Of course, the notion of a good society goes further. We might stress political rights and freedoms, intellectual and cultural development, stability of the family, social civility and tolerance, and so on. However, a high and widely accessible level of *material* well-being is probably a prerequisite for most other kinds of advancement, quite apart from being a worthy goal in itself.⁴ Economists and policy makers therefore do well (and have enough to do!) by concentrating on this aspect alone.

It's tempting to suggest that the state of material well-being of a nation is captured accurately by its *per capita gross domestic product*, GDP, which keeps track of domestically produced income per person in a country over a given year. (Or one might invoke its equivalent cousins, gross national income GNI and gross national product GNP, which add in net income from other countries: dividends, interest or repatriated profits.) Indeed, ever since national economic development was adopted as a conscious goal,⁵ there have been long phases during which development was judged exclusively by the yardstick of per capita income growth (of GDP or GNI or GNP). In the last few decades, this practice increasingly has come under fire from various quarters. The debate is longstanding, as the quotations at the beginning of this chapter suggest.

We must be careful here lest we be accused of setting up an obvious strawman. No one in their right mind would ever suggest that economic development be identified, in a *definitional* sense, with the level or growth of per capita income. It is perhaps universally accepted that development is not just about income, although income (economic wealth, more generally) has a great deal to do with it. For instance, we noted previously that economic advancement should not be restricted to a small minority. This means, in particular, that development is also the removal of poverty and undernutrition: it is lower inequality, it is an increase in life expectancy, it is access to sanitation, clean drinking water, and health services; it is the reduction of infant mortality; it is an opening door to knowledge and schooling.⁶ There is an veritable plethora of yardsticks. Paul Streeten's thoughts, summarized in the quotation at the beginning of this chapter, capture this "multidimensionality" very well.

Far more intriguing is the sharp focus of Robert Lucas' words (see quotation). At first they appear narrow, perhaps even missing the point, whereas the more holistic scenario sketched in the foregoing paragraphs seems pretty much the way to go. In thinking this we would be wrong. Neither Lucas nor any intelligent person believes that per capita income *is* development. What's hidden in these words is actually an

⁴This is not to suggest at all that it is *sufficient* for every kind of social advancement.

⁵For most poor countries, this starting point immediately followed World War II, when many such countries, previously under colonial rule, gained independence and formed national governments.

⁶The United Nations' *Millennium Development Goals* describe six targets: (1) eliminate extreme poverty and hunger, (2) universalize primary education, (3) achieve gender equality, (4) reduce infant and child mortality, (5) improve maternal health, (6) combat disease, such as HIV/AIDS and malaria, (7) build a "sustainable" environment, and (8) develop a global development partnership. You don't need to buy all these goals; the point is only to illustrate the implicit recognition of "development" as a multifaceted concept.

approach, not a definition. It is really a belief about the world, which is that *the universal features of economic development—health, life expectancy, literacy, and so on—follow in some natural way from the growth of per capita income*, perhaps with the passage of time. Implicit here is a belief in the power of aggregate economic forces to positively affect every other socioeconomic outcome that we want to associate with “development.” This outlook may be contrasted with the view that a correlation between per capita income and other desired features is not automatic, and that in many cases those connections may not be present at all. According to this latter view, per capita income fails as an adequate overall measure and must be supplemented by other indicators directly.

The debate implicit in the two quotations is not, therefore, about what development *means*, on which there is widespread agreement, though perhaps a bit too warm and fuzzy for precision. It is really about a world view — about the possibility of finding a smaller set of variables that correlates well with a multifaceted and complex process. Note well that, in a way, saying too much is saying too little. It may be that per capita income does not capture all aspects of development, but a ponderous assertion that no small set of variables ever captures the true nature of development and that there are always other considerations is not very helpful. In this sense, the view that economic development is ultimately fueled by per capita income may be taking things too far, but at least it has the virtue of attempting to reduce a larger set of issues to a smaller set, hopefully in a way that is supported by sound reasoning and empirical evidence.

This book implicitly contains a reduction as well, although not all the way to per capita income alone. In part, sheer considerations of space demand such a reduction. Moreover, we have to begin *somewhere*, so we concentrate implicitly on understanding two sets of connections throughout this book. One is how average *levels* of economic attainment influence development. To be sure, this must include an analysis of the forces that, in turn, cause average levels of income and other indicators to grow. The other connection is how the *distribution* of economic attainment, across the citizens of a nation or a region and across the nations of the world, influences development. The task of understanding these two broad interrelationships takes us on a long journey. In some chapters the relationships may be hidden in the details, but they are always there: levels and distribution as twin beacons to guide our inquiry.

In short, while we will study the many features of development is our ultimate goal, our approach to them lies through the two routes described in the previous paragraph.⁷

We begin with a summary of the historical experience of developing countries over the past few decades. We pay attention to per capita income; in particular to the way in which that income is distributed both across and within countries. We then discuss other indicators of development. Finally, we describe the structural characteristics of developing countries: the occupational distribution of the population and the shares of different sectors (such as agriculture, industry and services) in national income, and in this way set the stage for what is to come in this book.

⁷ Even this double emphasis might not be enough. For instance, the *Human Development Report* (United Nations Development Programme 1995) informs us that “the purpose of development is to enlarge all human choices, not just income. The concept of human development is much broader than the conventional theories of economic development.” Or, as Amartya Sen (1983) has warned, “Supplementing data on GNP per capita by income distributional information is quite inadequate to meet the challenge of development analysis.” There is much truth in these warnings, which are to be put side by side with Streeten and certainly contrasted with Lucas, but I hope to convince you that an understanding of our “narrower” issues will take us quite far.

2.2. Income and growth

2.2.1. Measurement issues. Low per capita incomes are an important feature of economic underdevelopment — perhaps *the* most important feature — and there is little doubt that the distribution of income across the world’s nations is extraordinarily skewed. Per capita incomes are, of course, expressed in takas, pesos, escudos, remimbi, and in the many other currencies of the world. To facilitate comparison, each country’s income (from all sources, but presumably predominantly in local currency) is converted into a common currency, typically U.S. dollars, and divided by that country’s population to arrive at a measure of per capita income. This conversion scheme is called the *exchange rate method*, because it uses the rates of exchange between the local and the common currencies to express incomes in a common unit.

The *World Development Indicators* of the World Bank contains such estimates for all countries. In 2024, the world produced approximately \$109.4 trillion of GNI for its 8.14 billion inhabitants, an annual average of around \$13,440 per person. Around \$8.3 trillion of this — 7.5% — came from all low-income and lower middle-income developing countries, accounting for 46% of the world’s population.⁸ Switzerland (with per capita income of around \$95,900) is 6 times as rich as Chile (\$15,750), over 7 times as rich as China (\$13,660), about 36 times as India (\$2650), and 150 times as rich as the Democratic Republic of Congo (\$640). Figure 2.1 contrasts per capita gross national income in different countries with the populations of these countries. The figure speaks for itself.




Figure 2.1. Population and Per Capita Income, 2024. Source: World Development Indicators (World Bank)

This book is not written with my heart on my sleeve, but the implied disparities are staggering. No amount of fine-tuning in measurement methods can get rid of the stark inequalities that we live with. Nevertheless, both for a better understanding of the *degree* of international variation that we are talking about and for the sake of more reliable analysis of these figures, it is best to recognize at the outset that these measures could provide biased estimates of what is actually out there.

(1) For one thing, underreporting of income is not uncommon in developing countries. Because tax collection systems are not as efficient as those prevailing in the industrialized market economies, there is a greater incentive to underreport income or

⁸For the 2026 fiscal year, the World Bank defined *low income* countries by those with 2014 per capita GNI under \$1,135. Many sub-Saharan African countries lie here, as do countries such as Afghanistan, with countries such as Myanmar, Tanzania, Rwanda and Uganda around the upper edge of that bracket. *Low middle-income* countries are those that lie between per capita incomes of \$1,135 and \$4,495; members of this group include India, Nicaragua, Nigeria, and Sri Lanka, with Iran or the Philippines around the upper edge. *Upper middle-income* countries include several of the richer Latin American economies, such as Argentina and Brazil, and countries such as Georgia, South Africa and Malaysia, with China and Bulgaria at the upper edge. They span the range between \$4,496–\$13,935. The *upper-income* countries make up the rest: the US, Canada, Western and Northern Europe, Australia, Japan, Singapore, and some Middle East countries.

output for tax purposes. The national accounts may not be comprehensive as well. In addition, the proportion of income that is actually generated for self-consumption is relatively high in developing countries; by subsistence farmers, for instance, who grow crops that they themselves consume. Such outputs may not be reported adequately. So these “shadow sub-economies” could come up to a large fraction of official national income, ranging from 30-50% in South and Southeast Asia (though with India probably at around 20-30%), 35-45% in Latin America, and climbing to 40-60% for sub-Saharan Africa; see Acharya (1981), Medina and Schneider (2019) and measures of informality from the World Bank Enterprise Surveys or the International Labor Organization.

(2) A more serious discrepancy arises from the fact that *prices* for many goods in all countries are simply not appropriately reflected in exchange rates. This is only natural for goods and services that are not internationally traded. Exchange rates are just prices, and the levels of these prices depends *only* on commodities (including capital) that cross international borders. The prices of nontraded goods, ranging from infrastructure to haircuts, do not affect exchange rates. What’s interesting is that there is a systematic way in which these nontraded prices are related to the level of development. Because poor countries are poor, you would expect them to have relatively low prices for nontraded goods: their lower incomes do not suffice to pull these prices up to international levels. This logic suggests that a conversion of all incomes to U.S. dollars using exchange rates *underestimates* the real incomes of poorer countries. This can be corrected, with incomes being recalculated to reflect what is known as *purchasing power parity*, or PPP. Many datasets, including the World Development Indicators, now also publish income data in this revised format.

Briefly (see box for more details), *international prices* are constructed for an enormous basket of goods and services by averaging the prices (expressed, say, in dollars) for each such good and service over all different countries. National income for a country is then estimated by valuing its outputs at *these* international prices. In this way, what is maintained, in some average sense, is *parity* in the *purchasing power* among different countries. Thus we call such estimates PPP estimates, where PPP stands for “purchasing power parity.”

PPP: The International Comparison Program

Here’s a little story: Asia’s weight in world output fell from 7.9% in 1985 to 7.2% in 1990 — and yet Asia was by far the fastest growing region during this period⁴. This same period also witnessed a sharp decline in some Asian exchange rates against the dollar. Now does that tell us something about the shortcomings of GDP exchange-rate estimates?

Actually, the trouble with market exchange rates for income calculations is not so much that they fluctuate, but that they do not fluctuate around the “right” average price, if “right” is to be measured by purchasing power. Even if exchange rates equalize the prices of internationally traded goods over time, substantial differences remain in the prices of nontraded goods and services such as housing and domestic transportation. There is a simple reason for this: because developing countries have relatively low incomes, you would expect non-traded goods to be cheaper. By assigning international prices to a basket of goods and then estimating incomes relative to those prices, we’re carrying out a comparison that maintains “purchasing power parity”; hence the term “PPP incomes”.

The International Comparison Program (ICP) began as a research project in 1968. Initially funded by the World Bank and the Ford Foundation, the ICP carried out comparison programs for 1970, 1973 and 1975 before its absorption into the United Nations Statistical Division (UNSD). The UNSD coordinated regional comparisons carried out by regional commissions. Research on the conceptual basis of such comparisons was initiated and continued by two economists at the University of Pennsylvania, Alan Heston and Robert Summers, both connected with the ICP since its inception. Their efforts resulted in the Penn World Tables (PWT; also called the Heston-Summers data set), with national accounts for a large set of countries dating from 1950. Entries were denominated in a set of "international" prices in a common currency, which drew on the ICP. Hence, international comparisons of income can be made both between countries and over time.^b

The ICP continues to provide price and expenditure comparison for an expanding set of countries, further rounds being completed in 1980, 1985, 1993, 2005, 2011 and 2017. The latest round (at the time of writing) is 2021 for 176 countries, conducted by the ICP Global office, now a permanent part of the Development Data group of the World Bank.

The first step of the ICP is simple: to compare prices for a well-defined commodity in (say) Bangladesh relative to the US. If it costs \$2 to buy a kilo of potatoes in the United States but 100 taka in Bangladesh, then the "potato exchange rate" between the taka and the dollar is 50. Indeed, life at the ICP would be blissful if all we consumed were the lowly potato, but of course, there are thousands of items, all with their own prices. Some of these are traded, and the overall exchange rate as we know it comes from these prices. Many items are not traded. The difficulties lie in aggregating from the lowest commodity heading (e.g., potato) all the way back up to overall income.

In particular, the ICP must gather detailed data on prices for thousands of items. These items are then classified numerous expenditure categories (broadly under the consumption, investment and government expenditure headings). By an averaging procedure following a method suggested by statistician R. C. Geary, the average price for each category is obtained relative to the price for that category in the U.S. In this way, numerous relative prices (or "price parities") are made available for each country and each category. Finally, categories must be aggregated up to national income.

The procedure is hard going, and issues of standardization severe.^c For instance, how do we aggregate different items as we go up to national income? Goods and services consumed in different countries are not all consumed in the same proportion. Their composition will vary depending on economic and social conditions, and any aggregation procedure needs to be sensitive to such compositional variations.^d There is an obvious truth in Deaton's (2010) pronouncement that "perhaps we are aiming too high when we try to construct a real income scale on which every country in the world can be placed," but an imperfect solution may be better than no attempt at all.

The end result is a PPP exchange rate for national output between two countries. If it is x taka to the dollar, it means that for every dollar spent on a bundle of goods and services in the United States, x taka would need to be spent on the "same" bundle goods in Bangladesh. The notion of sameness is based on the idea — at least in theory — that the "same" bundles must provide equivalent utility in the two countries.

There is no doubt that the ICP performs an enormously valuable function. At the same time, the maintenance of such a program hasn't been easy. For instance, it is not clear that the price information — for which one has to reply on many different domestic statistical agencies — is accurate. There are logistical problems in the collection and *coordinated* editing of the data, which requires close interaction between the national statistical offices and the regional offices of the relevant international organization. The World Bank, in its




Figure 2.2. Impact of PPP Adjustments for Share in Global Income. Source: International Comparison Program 2021 round; <https://www.worldbank.org/en/programs/icp>

role as global coordinator, has taken on several of these issues, by raising funds for the ICP as well as pushing for country-level participation in the Program.

It should also be noted that to the extent that PPP estimates affect income calculations, they have a bearing on formulae for aid allocations, and for this and related reasons, there might even be strategic reasons for “inaccuracy” (or at least procrastination) in the provision of price data. It could also have untoward and potentially manipulable implications for the construction of global poverty lines (Deaton 2010).

Finally, remember that while the data pioneered by Summers and Heston and the ICP are useful for “real” comparisons, exchange rate-based data are the appropriate ones to use for international financial transactions and capital flows.

^aSee *The Economist*, May 15, 1993.

^bThe PWT are housed at the Center for International Comparisons at the University of Pennsylvania; see <http://pwt.econ.upenn.edu/>. Apart from income data, the PWT also offers data on selected countries’ capital stocks and demographic statistics. In the revised GDP calculations based on PPP, Asia’s share in world output in 1990 jumped from 7 to 18%.

^cSee the World Bank’s site on the ICP for more detail: http://siteresources.worldbank.org/ICPEXT/Resources/ICP_2011.html.

^dAs an example, air travel accounts for a higher percentage of British consumption than it does in Cameroon. If we use the average of British and Cameroonian shares, air travel can have a large effect on Cameroon’s overall price level.

PPP estimates of per capita income go some way toward reducing the astonishing disparities in the world distribution of income, but certainly not all the way. Figure 2.2 shows how different regional and economic groups gain or lose when moving from exchange-rate measures to PPP measures. On the whole, richer regions lose, and poorer regions gain, in large part — as already discussed — due to the adjustments necessitated in the prices of non-traded goods. For a more granular account, consult Figure 2.3 which compares PPP and exchange-rate calculations of per-capita GNI. The setting is just the same as in Figure 2.1, with countries ranked in ascending order of exchange-rate GNI per-capita. The new, brighter line shows the PPP computations. As expected, the PPP numbers are relatively higher for poorer countries, because nontraded prices are lower in poorer countries, as we’ve discussed. World per-capita

GDP PPP in 2024 is \$24,083, well above the exchange rate estimate of \$13,670. And now the low and lower-middle income populations earn 17% of the world's income (as opposed to just 8.3% under the exchange rate calculations) — a welcome adjustment, but still dramatically low.

(3) There are other subtle problems. Income measurement, even when it accounts for the exchange-rate problem, uses market prices to compare apples and oranges; that is, to convert highly disparate goods into a common currency. The theoretical foundation for this is that market prices truly reflect preferences as well as relative scarcities. There are several objections to this argument. Not all markets are perfectly competitive; neither are all prices fully flexible. We have monopolies, oligopolistic competition, and public sector companies⁹ that sell at dictated prices. There is expenditure by the government on bureaucracy, on the military, or on space research, whose monetary value may not reflect the true value of these services to citizens. Moreover, conventional measures of GNI or GDP ignore costs that arise from externalities — the cost of associated pollution, environmental damage, resource depletion, or human suffering due to displacement caused by “development projects” such as mining, large-scale industry, dams, railways, and so forth. In all of these cases, the going prices do not capture the true marginal social value or cost of a good or a service.




Figure 2.3. PPP and exchange-rate GDP per capita, 2024.

All these problems can be mended, in principle, and sophisticated measures of national income do so to a large extent. Distortions in prices can be corrected for by imputing and using appropriate “shadow prices” that capture true marginal values and costs. There is a vast literature, both theoretical and empirical, that deals with the concepts and techniques needed to calculate shadow prices for commodities. An estimated “cost of pollution” is often deducted in some of the measures of *net* national income, at least in industrialized economies. Nevertheless, it is important to be aware of these additional problems. With that said, let us turn to a brief account of recent historical experience.

2.2.2. Historical Experience. Over the period 1990–2020, the richest 10% of the world’s nations averaged a per capita income (PPP) that started a bit over 5 times the world average, with that ratio rising and then dropping to well under 4 by 2024. The corresponding ratio for the poorest 10% of countries relative to the world average has held steady at around 0.1; see Table 2.1. A broad constancy of extreme inequality among the very richest and poorest countries has been maintained since 1960, with some improvement (as just discussed) in the first two decades of the 21st century. As Parente and Prescott (2000) quite correctly observe, interstate disparities *within* the United

⁹In many countries all over the Third World, sectors that are important or require bulk investment, such as iron and steel, cement, railways, and petroleum, are often in the hands of public sector enterprises.

States do not even come close to these international figures. In 2010, the richest state in the United States (not counting DC) was Alaska and the poorest was Mississippi, and the ratio of per capita incomes worked out to slightly over 2!

	1990	1995	2000	2005	2010	2015	2020	2024
top 10%/world average	5.18	5.30	5.37	5.16	4.56	4.22	4.01	3.87
bottom 10%/world average	0.12	0.10	0.10	0.10	0.11	0.11	0.11	0.10

Source: The World Bank, <http://data.worldbank.org/>.

Table 2.1. Richest and Poorest 10% Relative to World Average GDP Per Capita.

Of course, the fact that the richest and poorest 10% of countries bear approximately the same ratio of incomes (relative to the world mean and to each other) does not suggest that the *entire* world distribution of incomes has remained stationary. To get some idea of this, note first that world GDP per capita grew at an average annual rate of about 1.5% per year over 1970–2010, the rate picking up to 2.7% over 2011–2020.¹⁰ But East Asia danced to a tune all its own. Japan was initially the most visible, averaging 5.3% in per-capita GDP growth over 1960–1990, but similar or even higher growth rates over the same period were also characteristic of other East Asian economies: Korea (6.1%), Hong Kong (6.6%), Indonesia (3.8%), Malaysia (4.2%), Singapore (6.4%), or Thailand (5.1%). Despite the East Asian financial crisis in the late 1990s, these countries continued their economic advance into the 21st century, growing at annualized rates well in excess of 3% over 1990–2010, while Japan slowed to under the world average.¹¹

Impressive as these rates are, they are dwarfed by China’s phenomenal performance. Between 1980 and 1990, China’s per capita income grew at an annual rate of 7.6%. The corresponding figure for 1990–2010 is even higher: around 9.5%, with the rate slowing to what one can only describe as a *relatively* sedate pace of 5.9% over 2011–2020.

Other large economies in Asia have begun to move. India motored along at a modest rate of 2.6% per year over 1960–1990, which picked up to 3.6% over the 1990s, then to a remarkable 6.2% over 2000–2010 and a robust 4.3% over 2011–2020.

In contrast, much of Latin America languished during the 1980s. After a relatively high rate of economic expansion in the two preceding decades — at around 2.9% annually — growth slowed to a crawl. In many cases there was no growth at all. Over 1980–1990, during the so called “lost decade” for Latin America, per-capita income for the region *declined* by an average of over 0.7% year over year, leading to an overall drop of around 10%.¹² Down went Argentina (-2.9% *annualized* over 1980–1990), Brazil (-0.5%), Mexico (-0.3%), Peru (-3.0%), Uruguay (-0.7%) and their neighbors. Only Chile (2.1%) and Colombia (1.4%) had significantly higher per capita income in 1990 than they did in 1980. It is true that such figures should be treated cautiously, given the extreme problems of accurate GNP measurement in high-inflation countries, but they can’t be ignored either. With some notable exceptions (Chile, 4.7%, and Argentina, 3.6%), income growth remained slow in the 1990s; around the world average at a bit less

¹⁰ All growth numbers are taken from databank.worldbank.org unless otherwise indicated.

¹¹ Japan grew by less than 1% per year over this 1990–2010, while the numbers remained high for Korea (4.4%), Hong Kong (2.9%), Indonesia (3.3%), Malaysia (3.6%), Singapore (4.0%) and Thailand (3.4%).

¹² See Morley (1995) on the fortunes of Latin America over this period.

than 1.6%. We see a broader recovery over 2000–2010, with average growth rates well in excess of 2% annually; take note of Argentina (3.3%), Brazil (2.4%), Chile (2.6%), Peru (4.3%) and Uruguay (3.0%), though one of Latin America's largest economies, Mexico, has not fared quite so well (0.8%). Progress remains slow at 1.1% over 2011–2020, with countries such as Brazil and Argentina essentially stagnating during this time.

Similarly, much of Africa stagnated or declined over the 1980s. Sub-Saharan Africa as a whole declined at an annual per-capita rate of over 1%, and things were not any better in the 1990s (-0.4%), though the first decade of this century has seen a relative improvement (2.2%). Nigeria experienced substantial declines of per capita income through the 1980s (-1.6% per annum), and essentially stagnated through the 1990s, before pulling back to a more robust recovery over 2000–2010 (3.9%), slowing to 0.7% over 2011–2020. A similar pattern is visible for Tanzania in the 1980s and 1990s, with robust recovery rates of 4% (2000–2010) and 2.3% (2011–2020). Kenya barely grew in per capita terms in the 1980s, and continued to decline in the 1990s before recovering to some extent in 2000–2010; overall conditions (0.2%) over this *thirty-year* period are near-stagnant. But we do see a remarkable recovery in the last decade, as Kenya posted 6.2% over 2011–2020. Uganda stagnated over the 1980s (-0.1%) before picking up pace and making substantial progress over 1990–2010, growing at over 3.5% annually, though the pace has slowed to 0.5% over 2011–2020. A spectacular turnaround has been underway in Rwanda, crippled by negative growth in the 1980s (-1.2%) and 1990s (-0.7%) before a sustained recovery of well over 4.8% per year over 2000–2020. In dismal contrast, Burundi's negative growth rate of 3.2% in the 1990s is barely compensated for by its near-stagnation over 2000–2010 (0.4%) and slow growth over 2011–2020 (1.7%). One of the largest countries in Africa, the Democratic Republic of the Congo, went into veritable freefall over the 1980s (-2.2%) and the 1990s (-8.5% *annually*) before its turnaround of around 1.8% annually over 2000–2010 and recent sprint at 6.0% over 2011–2020.. And what of Zimbabwe, a country that stagnated in the 1980s (0.7%) and 1990s (-0.3%) before entering a freefall of its own (-4.8%) over 2000–2010?

We began this account with world incomes, and we end with another benchmark — the growth experience of country-members of the Organization for Economic Cooperation and Development (OECD). The twenty original members and the fourteen additions are largely developed countries, with a few middle-income members. Over 1970–1990, OECD growth was a bit over 2.4% annual, before falling to a more sedate 1.8% over the 1990s (but still exceeding the world average over this period) and then under the world average at 0.8% during 2000–2010, the slowdown marked by the Great Recession of 2007–2009. Over 2011–2020, growth has tracked the world average at around 2.6% per annum. The United States mirrors the OECD well, growing at over 2.2% over 1970–1990, falling slightly to a bit under 2.2% in 1990–2000, before slowing to 0.7% in 2000–2010 and then climbing to 2.7% in 2011–2020. Its overall growth is also uncannily close to that of the world average; see below.

This discussion isn't a comprehensive accounting of growth, but you can see that there's been a lot of churning in the international distribution of incomes. Such diversity demands explanation, but that demand is ambitious. No single story can account for the variety of historical experience. We know that in Latin America, the sovereign debt crisis triggered enormous economic hardship in the 1980s. In sub-Saharan Africa, low or negative growth is due in large measure to unstable government, civil strife and consequent infrastructural breakdown, as well as to high rates of population

increase. The heady successes of East Asia are not fully understood, but a conjunction of farsighted government intervention, a relatively equal domestic income distribution and a vigorous entry into international markets played an important role. We will take up these topics, and many others, in the chapters to come.

Growth experiences such as these can change the face of the world very quickly. One way to see this is to study the “doubling time” implicit in a given rate of growth; that is, the number of years it takes for income to double if it is growing at that rate. The calculation in this footnote¹³ reveals that a good approximation to the doubling time is 70 divided by the annual rate of growth expressed in percentage terms. So a country growing at 5% per year will *double* its per capita income every fourteen years! In contrast, a country growing at 1% per year will require seventy years. Percentage growth figures look like small numbers, but over time, they add up very fast indeed.

Thus it is quite possible for the world distribution of income to stay fairly constant in *relative* terms (and even at the extremes), while at the same time there is plenty of action *within* that distribution as countries climb and descend the ladder of relative achievement. Indeed, the several countries that we have cited as examples are no exceptions to the general picture of movement. To see this, we follow the convention of using the United States as a benchmark. As averages go, isn’t a bad benchmark. World average income was a bit under a quarter of US per-capita income in 1982; that ratio hasn’t budged at all. We express each country’s per capita GDP (PPP) relative to that of the United States, and look at how this ratio changes over time. Figure 2.4 shows the number of countries that experienced changes in income (relative to that of the United States) of different magnitudes over the years 1990–2020.¹⁴

Figure 2.4 indicates two things. First, over 70 countries — close to half — changed their position relative to the United States by an average of *one percentage point or more per year*, over the period 1990–2020. Second, there is a slight tendency for upward moves to dominate down-

ward moves, which partly accounts for the fact that you see an improvement in the world distribution taken as a whole, as corroborated by Table 2.1. (In the previous edition of this book, there was practically no net movement between 1960–1985.) This observation is cause for much hope and some trepidation: the former, because it tells us that there are probably no traps to ultimate economic success, and the latter, because it




Figure 2.4. Annual percentage change in PPP income of 158 countries relative to U.S. levels, 1990–2020. Source: The World Bank; <http://databank.worldbank.org>.

¹³A dollar invested at $r\%$ per year will grow to two dollars in T years, where T solves the equation $[1 + (r/100)]^T = 2$. This means that $T \ln_e [1 + (r/100)] = \ln_e 2$. However, $\ln_e 2$ is approximately 0.7, whereas for small values of x , $\ln_e (1 + x)$ is approximately x . Using this in the equation gets you the result.

¹⁴See Chart 10 in Parente and Prescott (1993) for an earlier version.

seems all too easy to slip and fall in the process. Economic development is probably more like a treacherous road, than a divided highway where only the privileged minority is destined to ever drive in the fast lane.

This last statement must be taken with some caution. Although there appears to be little evidence that very poor countries are doomed to eternal poverty, there is some indication that low incomes are very sticky. Even though we will have more to say about the hypothesis of ultimate convergence of all countries to a common standard of living (see Chapters 3–9), an illustration may be useful at this stage. Following Quah (1993), we can use per capita income data to construct “mobility matrices” for countries. First convert all per capita incomes to fractions of the *world’s* per capita income. Thus, if a country has per capita income of \$1,000 and the world average is \$3,000, assign the number 1/3. Now create categories that we will put each country into. Quah used the following categories (you can certainly construct others if you like): income less than a quarter of the world average (1), income between a quarter and half the world average (2), income between half the world average and the world average (3), income between the world average and twice the world average (4), and income exceeding twice the world average (5).

Now imagine doing this exercise for two points in time, with a view to finding out if a country transited from one category to another during this period. You will generate what we might call a *mobility matrix*. Figure 2.5 illustrates such a matrix for the period 1990–2020. On the very left, as you go down the rows, you can read off the number of countries in the various categories in 1990; e.g., 29 countries in the 143-country set had per-capita GDP less than a quarter of the world average in 1990. The columns house the corresponding categories in 2020. Each entry in each cell records the percentage of countries that made the transition from the row category (1990) to the column category (2020). For instance, 17% of the countries in category 1 made it to category 2 over the 30-year period. A matrix constructed in this way gives you a fairly good sense of how much mobility there is across nations. For instance, the presence of very high numbers on the main diagonal, consisting of those special cells with the same row and column categories, indicates low mobility. According to such a matrix, countries that start off in a particular category have a high probability of staying right there. Conversely, a matrix that has the same numbers in *every* entry (which must be 20 in our 5 × 5 case, given that the numbers must sum to 100 along each row) shows an extraordinarily high rate of mobility. Regardless of the starting point in 1990, such a matrix will give you equal odds of being in any of the categories in 2020.

Obs	Cat	①	②	③	④	⑤
29	①	72	17	10	0	0
29	②	34	45	17	3	0
30	③	3	13	60	23	0
25	④	0	0	40	52	8
30	⑤	0	0	0	10	90

Figure 2.5. The Income Mobility of 143 Countries, 1990–2020.

With these observations in mind, continue to stare at Figure 2.5. Notice that middle-income countries exhibit more movement than the poorest or richest countries. For instance, only 45% of the countries in category 2 remained where they were in 1990. In contrast, over 70% of the countries in category 1 remained where they were in 1990, and *none* of them made it over the world average by 2020. Likewise, 90% of the richest

countries in 1990 remained where they were relatively in 2020.¹⁵ This is interesting because it suggests that although everything is possible (in principle), a history of underdevelopment or extreme poverty puts countries at a tremendous disadvantage.

That's pretty trite, you might say: after all poverty should feed on itself and so should wealth, but on reflection you will see that this is really not so. There are many "leapfroggy" reasons to think that historically low levels of income may be *advantageous* to rapid growth. New technologies are available from the more developed countries. The capital stock is low relative to labor in poor countries, so the marginal product of capital could be high. One has, to some extent, the benefit of hindsight: it is possible to study the success stories and avoid policies that led to failures in the past. This account is not meant to suggest that the preceding empirical finding is inexplicable: it's just to say that an *a priori* guess does not yield straightforward answers. This "persistence versus convergence" theme is an important one that will recur through the book.

There is actually a bit more to Figure 2.5 than lack of mobility at the extremes. Look at the next-to-poorest category (those with incomes between one-quarter and one-half of the world average in 1990). Note that 17% of these countries went up by one category, but 45% of them *dropped* to the lowest category. It is not only the lowest-income countries that might be caught in a very difficult situation. There is the additional possibility that the middle of the distribution is like a slippery banana peel and a catapult all in one: it could take you in either direction quite easily.

To summarize, then:

(1) Over the period 1990–2020, the *relative* distribution of world income appears to have been quite stable. The richest 10% of nations averaged a per capita income around 3.8 times the world average; the poorest 10% had 13–14% of the same average. By any standards, this disparity is staggering, and especially so when we remember that we are talking about incomes that have been corrected for purchasing power parity.

(2) The fact that the overall distribution has remained reasonably stationary does *not* mean that there has been little movement within the world distribution. Of particular interest is the rise of East Asia and the languishing of other regions, sub-Saharan Africa and several countries in Latin America, followed in turn by spots of resurgence in the latter regions. Diverse growth experiences such as these can change the economic composition of the world in the space of a few decades. Nonetheless, a one-line explanation for this diversity remains elusive, which isn't surprising.

(3) The observation that several countries have changed relative positions suggests that there are no ultimate traps to development. At the same time, a history of wealth or poverty does seem to stick to some degree. Mobility appears to be highest somewhere in the middle of the wealth distribution, whereas a history of underdevelopment or extreme poverty appears to put countries at a serious disadvantage.

(4) That history might matter in this way is an observation that requires a careful explanation. Poor countries do seem to have some advantages. They can use, relatively free of charge, technologies that are developed by their richer counterparts. Scarce capital in these countries should display a higher rate of profit, because of the law of diminishing returns. They can learn from mistakes that their predecessors have made. In this way differences across countries should iron themselves out over the longer

¹⁵ Of course, our categories are quite coarse and this is not meant to suggest that there were no relative changes at all among these countries.


run. Thus the observation that history matters in maintaining persistent differences needs more of a justification than might be apparent at first glance.

2.3. Income Inequality Within Countries


The international disparity of national income is only one indication that something is fundamentally askew with global development. To this we must add the inequalities observable *within* the vast majority of developing countries. It is commonplace to see enormous wealth coexisting with great poverty, especially in the great cities of the developing world: in Lagos, Mumbai, Rio de Janeiro, Manila, or Mexico City. Not that inequalities don't exist in the developed world — they certainly do — but coupled with low average incomes, these disparities in developing countries result in a particularly visible explosion of poverty and destitution.

We will have much more to say on the topic of income distribution later in this book. However, it is useful to get an overall feel for the magnitude of the problem by looking at some data. Figure 2.6 summarizes information on inequality in 130 countries, with population above 1m, from the World Inequality Database, arranged in order of increasing per capita GDP. For each country, the figure records the income shares of the poorest 40% of the population as well as the income shares of the richest 20% of the population.

Look first at Panel (a) of Figure 2.6, and do so first without reliance on the fitted solid curves. (The Appendix contains this information for a specific list of countries.) The Figure plots the income shares of the richest 20% and the poorest 20% of the population in each country, represented by a vertically aligned pair of a circle and a triangle. The circles mark the shares of the richest 20% of the population within each country. They cluster between 60–70% for the poorer countries, falling — although with considerable scatter — into the 50-60% and then the 40-50% range as we move over the cross-section of poor to rich countries. There is a hint of a reversal in this trend at the very end, and the fitted curve appears to pick up this tendency. Now look at the heights of the triangles. In parallel vein, there is a slight upward movement,



(a) The full range of per-capita GDP



(b) Per-capita GDP below \$50,000

Figure 2.6. Income shares of poorest 40% and richest 20% in various countries. Source: The World Inequality Database and The World Bank; <http://databank.worldbank.org>.

certainly not as pronounced, in the income share of the poorest 40%, which appears to generally fluctuate in a range around 10%. Again, the fitted curve appears to suggest a slight reversal in this trend in the richest sub-range of countries.

Cross-sectional patterns aside, these shares represent sizable inequalities. When compounded with the inter-country differences that we've already discussed, we see that the poor in the developing world are twice cursed, once by the low per-capita incomes of the countries they inhabit, and then again by the reality of their low shares in the within-country income distribution. The world is a very unequal place indeed.

We've mentioned a declining trend for inequality overall as we move over the cross-section of countries. Panel (b) of Figure 2.6 zooms in on the relatively poorer countries below \$50,000 of GDP per capita (PPP) in 2023. (This is just Panel (a) stretched out over its first half.) This exercise reveals a different pattern that used to be even more pronounced a few decades ago. The fitted curves now show that there is some tendency for inequality to rise and then fall over as we move over the cross-section of countries. The pattern is subtle enough that it would ordinarily not be of any particular import, but it is historically connected to a famous assertion made by Simon Kuznets (1955, 1963): that inequality rises and then falls over the course of economic development.

Here is a summary of the Kuznets argument. At very low levels of living, a country is predominantly agrarian and in subsistence mode. For such countries, there are limits on the income share of the rich, as everyone is largely in agriculture. Now, as development proceeds, new sectors open up — most notably, an industrial sector is often born — and the economy begins a *structural transformation*. This transformation initially benefits some well-placed individuals and groups, who snap up jobs and opportunities in the new sector. As a consequence, the income shares of the top population deciles rise, and there is higher inequality. The share of the poorest groups tends to fall at the same time, although this does not necessarily mean that their income goes down in *absolute* terms. However, as the transformation continues, economic gains tend to be distributed more equally — poorer groups acquire relevant skills and education, and subsequently catch up in income share. The overall graph connecting per capita income and inequality then exhibits an upside-down U shape.

I mention this because this so-called “inverted-U hypothesis” was once remarkably influential, leading to a viewpoint that argued for inequality as a necessary but transient evil. And it is tied to the idea of structural transformation, a topic so important that we shall presently devote an entire chapter to it. Alas, in line with Panel (a), there is very little of an inverted-U that we see in the world today, in part because there is not just one structural transformation, but a mix of several great transformations: agriculture to industry, industry to services, informal to formal, manual to digital, culminating in the foment today as we pass through the AI structural transformation. But Kuznets could not have foreseen all these changes, and so while his idea of a *single* inverted-U was overly optimistic, we only know that with the benefit of hindsight.

Moreover, the inverted-U hypothesis contains a deeper notion that is robust and important: the core idea that economic development is an uneven process. Some sector takes off initially, enriching its inhabitants; later the others slowly catch up. Each transformation generates its accompanying waves of uneven and compensatory changes. That might generate an inverted-U in inequality, not just once (as in the the agriculture-industry transition that interested Kuznets), but several times in history.

But over and above this core point, there can be no inevitability about any supposed correlation between inequality and development. Countries that pursue the broad-based availability of infrastructure or resources such as health services and education, and policies that permit widespread and equal access to new opportunities, will in all likelihood find that economic growth is distributed relatively equally among its citizens. Countries that neglect these features will fare otherwise. Indeed (as you will be warned over and over again), variation in some variable *across* countries is not necessarily to be equated with a change in that variable as a *particular* country develops over time. The cross-sectional diagrams are interesting and sometimes useful, but we must refrain from taking them too seriously as guidelines to the fates of *individual* countries.

2.4. The Many Faces of Underdevelopment

The highly unequal distribution of incomes — and its variation across countries — suggests that excessive reliance on income per capita as a reliable indicator of overall development might be dangerous. A relatively prosperous country may fare poorly on some of the commonsense indicators of “human development”, such as literacy, access to drinking water, low rates of infant mortality and life expectancy. In fact, we also need to entertain the additional possibility that factors beyond just per capita income, or even the *distribution* of that income, may be at work. The social and economic empowerment of women may serve to significantly reduce infant mortality and (more generally) raise the health and nutritional status of children, yet neither income nor its equal distribution across households fully guarantees the empowerment of women. Likewise, a country that promotes popular science and health programs might be a welcome outlier in health or secondary education categories, even though income may be low or poorly distributed. Later in this section, we will emphasize the overall correlation of human development with per capita income, but it is always worthwhile to be sensitive to the outliers, because they tell their own story.

Consider Table 2.2. Each of the three tables here refers to a different country, and below the country name I have recorded the income shares of the poorest 40% and the richest 20% of its inhabitants in 1990. First compare Sri Lanka with Guatemala. Notice that Sri Lanka is *substantially* poorer in per capita terms than Guatemala. Yet Sri Lanka has a relatively equal income distribution; the share of the poorest 40% is 2.6

Sri Lanka (13, 57) in 1990	1990	2000	2007
Per capita GNI (PPP)	1450	2670	4210
< 5 Mortality (/1000)	29	21	21
Adult Literacy %	87	91	91
Life Expectancy (yrs)	69	71	74
Guatemala (5, 70) in 1990	1990	2000	2007
Per capita GNI (PPP)	2360	3470	4490
< 5 Mortality (/1000)	77	47	39
Adult Literacy %	46	69	73
Life Expectancy (yrs)	62	70	70
Pakistan (12, 57) in 1990	1990	2000	2007
Per capita GNI (PPP)	1260	1690	2540
< 5 Mortality (/1000)	130	108	90
Adult Literacy %	26	43	54
Life Expectancy (yrs)	61	64	66

Table 2.2. Some development indicators for Sri Lanka, Guatemala and Pakistan. Source: The World Bank, <http://databank.worldbank.org>, UNDP, *Human Development Report*, and World Inequality Database.

times as high than it is in Guatemala. Now look at some of the human development indicators for these two countries. Life expectancy was a good seven years higher in Sri Lanka in 1990. The difference has narrowed by 2007, but it's still there. Much of this difference came from the huge difference in child mortality rates, defined as the number of children (per thousand live births) who die before the age of 5. In Sri Lanka this figure was 29 per thousand in 1990; in Guatemala it was more than two and a half times higher. Even in 2007 it is close to twice as high. Finally, Sri Lanka had an adult literacy rate of 87% in 1990, climbing to over 90% by 2007. In Guatemala, that number was 46% in 1990 and 73% in 2007. Looking at these comparisons, it is hard to escape the conclusion that the highly unequal distribution of income in Guatemala is responsible, at least in part, for these differences in some natural yardsticks of development.

However, that isn't the whole story. Look at the case of Pakistan. The same human development indicators are nothing short of dreadful (though there are countries that do worse), and they are certainly significantly worse than their counterparts for Guatemala. Of course, one reason is that PPP per capita income in Pakistan is relatively low. But the income number is not entirely incomparable to that for Sri Lanka: certainly, 2007 Pakistan is in line with 2000 Sri Lanka from a narrow income perspective. The income distribution in Pakistan is also similar to that in Sri Lanka. Yet in 2007, Pakistan has human development indicators orders of magnitude worse than those for Sri Lanka in 2000. Clearly, government policies or cultural attitudes, such as those concerning education, health or the treatment of women play significant roles.

The Human Development Index

Many of the *direct* symptoms of underdevelopment are easily observable and independently measurable. Undernutrition, disease, illiteracy — these are among the stark and fundamental ills that a nation would like to remove through its development efforts. For quite some time now, international agencies (such as the United Nations) and national statistical surveys have been collecting data on the incidence of malnutrition, life expectancy at birth, infant mortality rates, adult literacy rates, and various other direct indicators of the health, educational, and nutritional status of different populations.

As we've seen, a country's performance in terms of income per capita might be significantly different from the story told by these basic indicators. Some countries, comfortably placed in the "middle-income" bracket, nevertheless display literacy rates that barely exceed 50%, infant mortality rates close to or exceeding one hundred deaths per thousand, and undernourishment among a significant proportion of the population. On the other hand, there are instances of countries with low and modestly growing incomes, that have shown dramatic improvements in those basic indicators. In some categories, levels comparable to those in the industrialized nations have been reached.

The United Nations Development Programme (UNDP) has published the *Human Development Report* since 1990. One objective of this *Report* is to coalesce some of the indicators that we have been discussing into a single *human development index* (HDI). This is not the first index that has tried to put various socioeconomic indicators together. A forerunner is Morris D. Morris's "physical quality of life index" (Morris 1979), which created a composite index from three indicators of development: infant mortality, literacy, and life expectancy conditional on reaching the age of 1.

The HDI, as revised in 2010, begins with three components or dimensions. A "dimension index" is constructed for each component, and these are subsequently aggregated.

The first component is life expectancy at birth (this will indirectly reflect infant and child mortality). The second component is a measure of educational attainment. This measure is itself a composite: it combines two subindices, one tracking mean years of schooling, the other expected years of schooling viewed from the current year onwards. The last component is per capita income: the logarithm of this income is chosen, presumably on the grounds that there is diminishing marginal utility to higher incomes. In each case, a dimension index is composed by setting maxima and minima to the dimensions, and then tracking “achievement” by the ratio of the path already traversed along that dimension to the entire length of the path.^a The HDI is obtained by taking a symmetric geometric mean of the three dimension indices.

The creation of composites from such fundamentally different indicators as life expectancy and literacy is a bit like adding apples and oranges. Perhaps, rather than composites, the reader should view the different components and then judge the overall situation for herself. The advantage of a composite index is its visual simplicity and, of course, its political power: in this era of sound bites, it is far easier and appears to be more “scientific” to say that country *X* has an “index” of 0.8 out of 1, rather than to laboriously detail that country’s achievements (or failures) in five different spheres of development. The HDI might *look* scientific and the formulae used to create the final average might *look* intricate, but that is no reason to accept the implicit weighting scheme that it uses, because it is just as ad hoc as any other. Nevertheless, the HDI is one way to combine important development indicators, and for this reason it merits our attention.

The HDI creates, for each country, a final number somewhere between 0 and 1. The number is to be (tentatively) interpreted as the “fraction of ultimate development” achieved by the country in question. Because a notion of “ultimate bliss” is embodied in the HDI, and especially because the logarithmic transform of income is chosen, it is not at all surprising that the indicator is relatively varied among the poorer countries, but then flattens out sharply as we move into richer countries. This creates an artificial tendency for the index to more easily converge across countries than incomes would (low-indicator countries experience higher growth in their indicators).^b

At the same time, the *rankings* generated by the index are of interest because they illustrate how a relatively high-income country might fare so badly in meeting basic socioeconomic goals that its HDI lags behind that of a relatively poor country. One way to describe how this happens is to compare the difference in the two country rankings induced by these two measures. What about the examples in the main text: Sri Lanka, Guatemala, and Pakistan? The ranking approach justifies what we already saw on the basis of specific indicators. Sri Lanka has a positive rank differential of +10. Guatemala and Pakistan have negative rank differentials of -13 and -4, respectively.

There are all sorts of variations on the HDI. For instance, one can adjust each of the components for within-country inequalities. Or one might specialize the indicator (or more accurately, use the same aggregation approach) to other outcomes of interest, such as gender or poverty. UNDP’s website at <http://hdr.undp.org/en/> contains descriptions of these variations, as well as a comprehensive and easy-to-use database.

^aFor instance, the maximum is 83.2 years, presumably as of 2010, while the minimum is set to 20, so that a country with a life expectancy of x enjoys the dimension index $(x - 20)/(83.2 - 20)$.

^bIn particular, be careful of statements such as those made in the *Human Development Report* (1995, page 19): “the HDI of industrial countries (0.916) is *only* 1.6 times higher than that of developing countries (0.570), even though their real GDP per capita (PPP\$) is 6 times higher” (emphasis added). They have little meaning, given the use of log income as well as the notion of bliss points in the other dimensions. It’s like saying that an earthquake measuring 9 on the Richter scale is “only 12.5% more powerful” than one measuring 8!

There is little doubt, then, that per capita income, or even the equality of its distribution, does not serve as a unilateral guarantee of success in “human development.” This sentiment is captured very well in one of the views of development with which we started this chapter. At the same time, the apparently narrow perspective of mainstream economists, with its hard-nosed focus on per capita income as a summary statistic of development, may not be too out of line. It is arguable that although taking a wider and multidimensional view of development is *conceptually* correct, per capita GDP still acts as a fairly good *proxy* for most aspects of development.¹⁶ For instance, it can be argued that rising income levels generally translate into better health, nutritional, and educational standards in a population. Paying due attention to the exceptions does not mean that the general rule should be ignored. As you will see — and I don’t mind at all — I am oscillating between the statistical and the case-by-case perspective. I believe that this schizophrenic split can be healthy in moderation.




Figure 2.7. HDI versus GNI. Source: UNDP, *Human Development Report* and World Development Indicators (World Bank).

Figure 2.7 illustrates the argument for the Human Development Index (see box for a detailed description). Recall that the HDI combines different measures of development into a single index, and that country rankings according to HDI do not necessarily correspond to rankings according to income per capita. And yet the correlation between per capita income and HDI is very strong. Figure 2.7 plots per capita income (PPP) against the HDI for countries in 2023. Yes, the resulting plot is a scatter, but it closely hugs a logarithmic line of fit. Of course, log income is one of the components of the HDI, but it is only 1/3 of the components, and anyway the whole point of the *Human Development Report* has to been to argue that HDI goes well beyond per capita income.




Figure 2.8. Income Per Capita and Life Expectancy. Source: World Development Indicators (World Bank).

We don’t have to restrict ourselves to HDI. We could choose any indicator that interests us, and examine its correlation with per-capita income (or indeed, with any other variable we like). In this section, we illustrate the argument with three indicators of development: life expectancy at birth, the infant mortality rate, and the adult literacy rate. Of course, these indicators are not entirely independent of each other. For

¹⁶For additional information see the contributions of Anand and Harris (1994), Aturupane, Glewwe and Isenman (1994), Desai (1991), Naqvi (1995), Srinivasan (1994), and Streeten (1994).

instance, life expectancy *includes* the possibility of dying before the age of 1, which is infant mortality. Nevertheless, these are common indicators that enter into indices of development, such as the HDI or the physical quality of life index. Figure 2.8 plots the per capita income and life expectancy. The relationship is concave, of course (one cannot expect life expectancy to grow linearly with income throughout), but it is very tightly estimated. Figure 2.9 does the same for per capita income and infant mortality, and Figure 2.10 does so for per capita income and adult literacy.

The relationships between per capita income *alone* and of these variables are strikingly strong.¹⁷ The figures express the idea that per capita income is a powerful correlate of development, no matter how broadly we conceive of it.¹⁸

At the same time, I must qualify these observations. The diagrams do not convey causality: i.e., that per capita income directly *impels* change in some other feature of development. We're looking at correlations, and in a pretty loose way at that. We will return to this issue of causality at several points in the book. But there is even more to be said by way of qualification. We haven't even established that *within a country*, per capita income is correlated with social outcomes, causally or not. With country fixed effects properly accounted for in panel data, the evidence is mixed: see Easterly (1997). To summarize, we do not claim that simple cross-country studies can settle the issue conclusively, and we certainly do not propose that income is a complete determinant for all other facets of development.

To continue this delicate balancing act of hedging our interpretations, note moreover that even the correlation between per capita income and other indicators is strong but




Figure 2.9. Income Per Capita and Infant Mortality. Source: World Development Indicators (World Bank).




Figure 2.10. Income Per Capita and Adult Literacy. Source: World Development Indicators (World Bank).

¹⁷ Contributions arguing that higher per capita income is correlated with quality-of-life indicators include Mauro (1993), Pritchett and Summers (1995), Boone (1996) and Barro (1996).

¹⁸ Indeed, by looking at the actual *levels* of achievement in each of these indicators, rather than the ranking across countries that they induce, I have actually made life more difficult for the argument in favor of per capita income. In an influential book, Dasgupta (1993) showed that per capita income is correlated even more highly with other indicators of development if we use *ranks* rather than *cardinal measures*. In other words, if we rank countries according to their per capita GDP levels and redo this based on some other index (such as adult literacy, child mortality, etc.), then we find a high degree of statistical correspondence between the two sets of ranks. Because I have already carried out cardinal comparisons, I will skip a detailed discussion of these matters and simply refer you to Dasgupta's study for a more thorough reading.

far from perfect (otherwise the data would all lie on some smooth curve linking the two sets of variables). The imperfect nature of the relationship is just a macro-reflection of what we saw earlier with countries such as Sri Lanka, Pakistan, and Guatemala. Inclusion of the *distribution* of per capita income would add to and improve this fit, but even then matters would remain undecided: social and cultural attitudes, government policy, and the public demands for such policies, all would continue to play their role in shaping the complex shell of economic development.

In conclusion, the point of this section is not to discredit human development, but only to show that we must not necessarily swing our opinions to the other extreme and disregard per capita income altogether. It is in this spirit that we can appreciate the seemingly narrow quotation from Robert Lucas at the beginning of the chapter, without necessarily buying into the idea that economic growth *defines* development. In short, we must (and will) begin with a study of how per capita incomes evolve in countries. That said, it is only natural that we then continue to move to other pressing matters, such as the study of income distribution and the operation of various markets and institutions.

2.5. Structural Features

Our final objective in this chapter is to review some structural features of developing countries. We will examine these characteristics in detail later in the book.

2.5.1. Demographics. Very poor countries are characterized by both high birth rates, as well as high death rates in every age group. The overall death rate is not as high, though, because a high-birth-rate population is generally younger. As development proceeds and health interventions become widespread, death rates plummet downward. But, for a variety of reasons to be discussed later, birth rates remain high before they follow the death rates on their downward course. In the process, a gap opens up (albeit temporarily) between birth and death rates. This could lead to high population growth in developing countries, until the birth rate comes down as part of a “demographic transition” that dampens net population growth. Chapter 14 discusses the transition and related issues in detail.

Figure 2.11 summarizes some of these characteristics. Panel (a) plots both birth rates and death rates, customarily expressed per 1000, as we move over the cross-section of per capita incomes. Birth rates clearly show a decline over that cross-section. Death




Figure 2.11. Demographic Trends over a Cross-Section of Countries. Source: The World Bank; <http://databank.worldbank.org>.

rates display a more ambiguous pattern because of the younger age structure. Net population growth is the difference between birth and death rates but also adds in-migration (and subtracts out-migration). Panel (b) plots these numbers. As you can see, there is some tendency for migration from poor to rich countries to possibly outweigh the declining trend caused by subtracting death rates from birth rates alone.

High population growth has two effects. It means that overall income must grow faster to keep per capita incomes steady. At the same time, population growth helps income to grow, because there is a greater supply of productive labor. However, it is unclear who wins this seesaw contest: the higher output thereby produced or the larger population that drags per capita output downwards. The negative population effect may well end up dominant, especially if the economy in question is not endowed with large quantities of capital (physical or human).

A second effect of high population growth (or high birth rates, to be precise) is that the overall population is young, as we've already remarked upon. With high birth rates, a proportionately larger number of children are always entering the society, so the overall population is heavily weighted in favor of children. This may be quite delightful, as any of us who has grown up with several siblings and cousins knows, but it does not change the grim reality of economic dependence for those in poverty. There are many untoward consequences of a young population, and these include poverty, child labor, and low education.

2.5.2. Agrarian Share. Agriculture accounts for a significant fraction of production in developing countries. Indeed, given that substantial agricultural output is produced for self-consumption and so may not be picked up in the data, the proportion is probably higher than that revealed by the published numbers. Panel (a) of Figure 2.12 shows that the share of agriculture in overall GDP varies from

around 40% for the very poorest countries, through around 20% for middle income countries, all the way down to the low single digits (around 1-7% for the richest countries). The relationship is not linear but the connection is very tight indeed. The declining pattern exhibited by the share of the labor force in agriculture is even more striking. Panel (b) of Figure 2.12 illustrates. The share declines over the cross-section




Figure 2.12. Share of Output and Labor in Agriculture. Source: World Development Indicators (World Bank).

of countries, falling precipitously from a high of 80% to around 20-25% for middle-income countries, before declining to a number close to 0 for high-income countries.

There are two patterns here. First, the overall numbers for production and occupational structure suggest that agriculture often has lower productivity than other economic activities. This is not surprising. In many developing countries, capital intensity in agriculture is at a bare minimum, and there is often intense pressure on the land. Add to this the fact that agriculture, especially when not protected by assured irrigation and ready availability of fertilizer and pesticides, can be a singularly risky venture. Second, and precisely because land productivity — via mechanization and technical knowhow — is higher in richer countries, the labor share decline in agriculture is sharper than the output share decline. The latter exhibits a flatter pattern as technology compensates for the withdrawal of labor.

This cross-sectional pattern faithfully mirrors trends over time for a majority of countries. Most countries were historically agrarian, and then transitioned to non-agriculture. This is the leading instance of a *structural transformation*, a notion that we’ve already encountered in Section 2.3, and that we will study in detail in Chapter 6. That said, agricultural activity *still* forms a significant part of the lives of people living in developing countries. We will therefore devote a good part of this book to economic arrangements that bear on the rural sector: the hiring of seasonal labor, the leasing of land, and the operation of informal credit markets, though of course our discussions and models will have implications for other sectors as well.

2.5.3. Rural-Urban Migration. Given that developing countries are fully into the agrarian transformation, it is hardly surprising that an enormous amount of labor moves from rural to urban areas. Such large migrations deserve careful study. They are an outcome of both the “push” from agriculture, because of poverty and landlessness, and the perceived “pull” of the urban sector. The pulls are reinforced by a variety of factors, ranging from the comparatively high wages and worker protections offered in the organized urban sectors to the effect of the media in promoting the urban lifestyle as a desirable end in itself. The media is often misleading and so are the benefits of the organized sector, which are often accessible only to a lucky minority of workers.

Consider the rates of growth of the urban sector in developing countries. For the low-income countries as defined for 2024 by the World Bank, the average *annual* rate of urban population growth over the period 1960–2024 was 4.13% per year. Compare this with an average rate of population growth of 2.65% per year for the same countries over the same period of time. For the countries classified as low-middle income by the Bank, the corresponding numbers were 3.37% and 2.09%, while for the upper-middle income countries, they were 3.00% and 1.40%. Imagine, then, the pressure on the cities of these countries. This is not to say that such migrations are somehow undesirable. After all, now-developed countries went through these transitions as well. The fact of the matter, however, is that

	1964	1984	2004	2024
World	35.3	40.8	48.5	57.7
Low	16.3	24.1	30.9	38.1
Lower Middle	19.2	26.1	32.4	41.5
Upper Middle	27.5	35.7	51.9	69.3
High Income	63.7	71.9	77.0	81.4

Table 2.3. Urban Population Share in Different Income Groupings. Source: World Development Indicators (World Bank).

all these processes are accelerated in modern-day developing countries, and the speed-up imposes enormous strains. Table 2.3 summarizes this discussion by showing how the urban population share has evolved in different income groupings over 1960–2024. In 2024, over half the world’s population was urban.

2.5.4. Informality and Services. One piece of evidence that reveals these strains is the unusually large fraction of the population in developing countries that’s classified as being in the tertiary or “services” sector. Before looking at the data, it is useful to conceptualize matters a bit. Think of what we consume as our income increases. Our first needs are for food and clothing. As we have more income to spare we switch to industrial products: radio, television, bicycles, automobiles, and the like. At a still higher level of income we begin to register a high demand for services: banking, tourism, restaurants, and travel. It is not surprising, then, that developed countries allocate a large fraction of their nonagricultural labor force to the services sector. Countries such as Australia, the United States, the United Kingdom, Norway, and Sweden have about 80% of their nonagricultural total labor force in the services sector; the corresponding figures for other developed countries such as Japan are somewhat lower. That isn’t odd at all. What *is* odd is that many developing countries exhibit large fractions of the labor force in “services” as well! Table 2.4 shows that three quarters of the non-agricultural labor force in *low income countries* in 2024 are classified by the International Labor Organization as being in services as well! (The numbers for middle-income countries are lower, but still pretty high too.)

	2000	2023
World	65.6	68.0
Low Income	69.2	74.5
Lower Middle Income	65.8	62.9
Upper Middle Income	60.9	65.0
High Income	70.6	76.7

Table 2.4. Share of Non-Agricultural Labor in Services for Different Income Groupings. Source: International Labor Organization Modeled Estimates and World Development Indicators (World Bank).

Figure 2.13 makes the general point very clear. Expressed as a fraction of the nonagricultural labor force, the proportion in the services sector is not at all different from what we see in developed countries. At the same time, the proportion of people in agriculture *does* vary a great deal, as we have already seen. What we are seeing, then, for developing countries, is a classification of a large part of the labor force into “services” simply because such services are waiting positions or fall-back options for laborers lacking an industrial job. That is, the enormous services sector in developing countries is symptomatic of the development of the unorganized or *informal* sector, on which we will have more to say in Chapter 6, Section 6.9. This sector is the home of




Figure 2.13. Share of Services in Non-Agricultural Employment. Source: ILO Modeled Estimates and World Development Indicators (World Bank).

last resort — the shelter for the millions of migrants who have made their way to the cities from the rural sector. Micro service businesses, petty retail, small-scale trade and informal finance: all of it gets lumped under the broad rubric of services because there is no other appropriate category. The large size of this sector in developing countries is, in the main, a reflection of the inability of industry in these countries to keep up with the extraordinary pace of rural–urban migration. So if you think that the flip size of “rural” means “urban” in the sense of having a job in the government, or a registered factory or a commercial bank, think again.

2.5.5. International Trade. By and large, all countries, rich and poor, are significantly involved in international trade. Define a measure of openness to trade by

$$\text{Trade Openness} \equiv \frac{1}{2} \frac{\text{Exports} + \text{Imports}}{\text{GDP}} * 100.$$

A quick plot of trade openness against per capita GDP reveals that this ratio has some tendency to rise with per capita income, but there are significant variations. At the end of the day, trade openness has more to do with country size than per capita income. For instance, the United States is one of the least open countries in the world, with a trade openness of just 12%. But so are many other large countries, among them Brazil (17%), India (22%) and even China (18%)! Then again, there are small countries or regions such as Ireland (119%), Singapore (163%) and Luxembourg (202%) for which these ratios attain astronomical heights. The average (unweighted) value of trade openness across countries is around 45%. So trade is undoubtedly an important component of the world economy for both developed and developing countries — but the largest countries are largely closed, and the smallest countries are largely open.

The differences between developing and developed countries are somewhat more pronounced when we look at the *composition* of trade. Developing countries are often exporters of primary products: food and forestry items, raw materials, and cash crops. Countries such as Brazil, Nigeria, and Ghana come to mind as typical examples. In contrast, the bulk of exports from developed countries is in the category of manufactured goods, ranging from capital goods to consumer durables.

Of course, there are many exceptions to these broad generalizations. Countries such as Bangladesh, Vietnam, India, the Philippines, Mexico and of course China do have a very high proportion of manufactured goods in their exports, ranging from textiles to pharmaceuticals to semiconductors or automobiles.¹⁹ These countries and many of their compatriots are deliberately attempting to diversify their exports away from primary products. On the other hand, there are developed countries that export primaries (oil, natural gas, dairy, meat) to a great degree. Australia, New Zealand, and Norway are among them.




Figure 2.14. Share of Primary Exports in Total Exports. Source: World Development Indicators (World Bank).

¹⁹Exporters of electronics (Philippines, Vietnam, Malaysia) often re-export imported intermediates.

Figure 2.14 summarizes the overall trends by plotting the share of primary exports in total exports, against per capita income. It is clear that, on the whole, developing countries do rely on primary product exports, whereas the opposite is true for the developed countries. One explanation for this structure of international trade comes from the theory of *comparative advantage*, which states that countries specialize in the export of commodities in which they have a relative cost advantage in production. These cost advantages might stem from differences in technology, domestic consumption profiles, or the endowment of inputs that are particularly conducive to the production of certain commodities. We review this theory in Chapter 27. Because developing countries have a relative abundance of labor and a relative abundance of unskilled labor within the labor category, the theory indeed predicts that such countries will export commodities that intensively use unskilled labor in production. To a large extent, we can understand the aforementioned trade patterns using this theory.

At the same time, the emphasis on primary exports may be detrimental to the development of these countries. Primary products are particularly subject to large fluctuations in world prices, and this creates instability in export earnings for countries that depend heavily on such products. Additionally, over the longer run, as (consumer) primary products become less important in the consumption baskets of people the world over, a declining price trend might be evident for such products as well. In general, then, activities that have comparative advantage today might not be well suited for export tomorrow. The adjustment to a different mix of exports then becomes a major concern. Finally, technology often is assimilated through the act of production. If production and exports are largely limited to primary products, the flow of technology to developing countries may be affected. We discuss these issues in Chapter 28.

The *import* mix of developing countries is more similar to that of developed countries. Exporters of primary products often need to import primary products as well: thus India might be a major importer of oil and Mexico a major importer of cereals. Primary exports for each country are often concentrated in a handful of products, and there is no contradiction in the fact that primaries are both exported and imported. A similar argument establishes that although developed countries might export manufactured items, there is always a need for other manufactures that are in relatively short supply. The overall trade patterns for imports are depicted in Figure 2.15.




Figure 2.15. Share of Primary Imports in Total Imports. Source: World Development Indicators (World Bank).

We summarize: developing countries are likely to have a high ratio of primary goods in their total exports, but as far as imports are concerned, matters are not inverted: developing countries have a larger share of primary imports in total imports as well.

2.6. Summary

We began with a discussion of what the term *economic development* might mean. It is a multifaceted concept, embodying not just income and growth, but also achievements on other fronts: reductions in infant mortality, higher life expectancy, advances in literacy, widespread access to medical and health services, and so on. Per capita income is sometimes used as an (incomplete) indicator for overall economic development, but should not be identified conceptually with development in the broader sense.

We turned next to per capita income data for countries. Using exchange rates to convert local currencies into dollars, we obtained per capita income evaluated according to the *exchange rate method*. The disparities across countries is enormous. Some of this disparity is due to underreporting of income, but a far more serious problem arises from the fact that price levels are systematically different across countries: dollar prices for nontraded goods and services tend to be lower in developing countries. The *purchasing power parity method* attempts to correct for this by constructing international prices that are used to estimate national incomes. Cross-country disparities in per capita income are then smaller, but still large: the richest 10% of the world's nations averaged a per capita income that ran well over 40 times the corresponding figure for the poorest 10%, over the period 1990-2024.

There have been substantial changes in incomes for many countries. The meteoric rise of East Asia is a case to be noted. This case is contrasted with the fact that much of Latin America and sub-Saharan Africa languished (relatively speaking and often in absolute terms) in the 1980s and well into the 21st century. Thus, although the world distribution of income remained fairly unchanged in relative terms, there was plenty of movement within that distribution. However, there is evidence that a history of underdevelopment or extreme poverty feeds on itself. Using *mobility matrices*, we noted that middle-income countries have significantly higher mobility than either the poorest or the richest countries.

Next, we studied income distribution *within* countries. By and large, income is more unequally distributed in developing countries than in their developed counterparts, which suggests that the poor in developing countries are twice hit: once at the level of distribution across countries and then at the level of distribution within countries. We discussed patterns of within-country income distribution across countries at various levels of per-capita GDP. We took particular note of the famous Kuznets inverted-U hypothesis, which states that inequality rises and then falls over the course of economic development. While there is no particular evidence for that hypothesis today, we note for possible future use that the concepts of “uneven” and “compensatory” change buried within it could well be useful.

With income and income distribution out of the way, we then returned to the broader notion of development. The *Human Development Index* is the name given to a set of indicators developed by the United Nations Development Programme. It combines three indicators—life expectancy at birth, educational attainment, and per capita income—with weights to arrive at a combined index. We noted that just because an overall index is provided does not mean it should be necessarily taken seriously: the weights are, of course, quite arbitrary. Nevertheless, the overall idea of human development is a laudable attempt to conceptually go beyond per capita income as an operational measure of development.

Nevertheless, per capita income isn't a bad predictor of human development. We showed that the correlations between per capita income and other variables that describe "human development" are high, even if attention is restricted only to the subsample of developing countries.

Finally, we described some structural characteristics of developing countries. We looked at *demographic characteristics* and showed that there is a general tendency for population growth rates to decline with increased per capita income. We discussed very briefly some of the effects of population growth on per capita income. We then studied *occupational and production structure*: agricultural activity accounts for a significant fraction of occupations in developing countries. At the same time, the rates of *rural-urban migration* are very high indeed. In part, this is reflected in the observation that a large fraction of the nonrural labor force is engaged in a nebulous activity called "services." This category includes all sorts of informal activities with low setup costs, and in developing countries is a good indicator of urban overcrowding. At the end, we discussed *patterns of international trade*. Developing countries are largely exporters of *primary products*, although this pattern shows change for middle-income countries. Primary product exports can be explained using the theory of *comparative advantage*, although we note that primary product exports have intrinsic problems, such as a strong tendency for their international prices to fluctuate, which creates instability in export revenues. The import mix of developing countries is, however, also more primary-oriented, which creates an additional source of uncertainty as well.

Appendix: Inequality and Per-Capita Incomes Within Countries

We record here the shares of the poorest 40% and the richest 20% of income groups in a subset of countries arranged from poorest to richest. Notice that inequality is generally much higher for the poorer countries than the richer countries. Generally, the shares of the bottom 40% do rise over the cross-section of countries arranged by income, and the shares of the top 20% falls. But we need to qualify this by noting as we move from the poorest into the middle group of countries, there is also some tendency for these shares to go the other way — in the direction of rising inequality.

Reading across and down the Table is an instructive experience. At first glance, we see vestiges of Kuznet's famous inverted U over the cross-section of countries. The shares of the poorest 40% tend to fall and rise again as we move from poor to middle income countries and then on to rich countries, while the opposite is true of the income shares of the richest 20%. But it should also be noted that the majority of the countries in the middle that exhibit this high inequality is from Latin America. So: is this Kuznets at work, or something peculiar to Latin America?

This kind of question brings out the inherent limitations of a cross-sectional approach when what we are really after is the experience of *individual* countries over time. There is no reason to believe that the fortunes of any one individual country are going to mirror the cross-sectional pattern. If we study individual countries over time, however, there is additional evidence that inequality is on the rise in many of them. We shall do so in Chapter 11.

0–10,000				10,000–25,000			
		Income Shares					
Country	GDP pc	Bot. 40%	Top 20%	Country	GDP pc	Bot. 40%	Top 20%
Somalia	1,557	10	62	India	10,323	11	67
Malawi	1,830	10	62	Philippines	10,986	10	60
Afghanistan	2,202	14	54	Sri Lanka	14,456	11	62
Sierra Leone	3,371	11	60	South Africa	15,194	4	79
Ethiopia	3,061	12	59	Peru	17,011	5	71
Nepal	5,395	13	55	Egypt	18,525	13	61
Pakistan	6,037	14	56	Colombia	20,944	3	76
Nigeria	6,207	11	57	Brazil	21,176	5	73
Bangladesh	9,148	13	56	Thailand	23,519	7	67
Morocco	9,843	10	62	Mexico	24,855	4	74
25,000–40,000				40,000+			
		Income Shares				Income Shares	
Country	GDP pc	Bot. 40%	Top 20%	Country	GDP pc	Bot. 40%	Top 20%
China	25,179	8	58	Korea	52,203	13	47
Costa Rica	28,075	4	67	Spain	53,230	15	48
Argentina	30,082	10	58	U.K.	57,915	13	51
Chile	32,801	5	71	France	58,318	13	49
Uruguay	34,471	10	61	Sweden	67,259	17	45
Malaysia	36,417	13	54	Germany	68,693	13	51
Turkey	42,326	10	67	Australia	70,513	10	50
Croatia	45,596	13	49	U.S.	82,305	9	62
Portugal	47,426	13	51	Norway	100,437	16	45
Japan	49,897	11	54	Singapore	143,786	13	58

Table 2.5. Income Shares of Poorest 40% and Richest 20%, Selected Countries with Varying Per Capita PPP GDP, 2023. Source: The World Inequality Database and The World Bank; <http://databank.worldbank.org>.