

Interrogating the holy grail of productivity growth

Jayati Ghosh¹ [University of Massachusetts at Amherst, USA]

Copyright: Jayati Ghosh, 2021

You may post comments on this paper at

<https://rwer.wordpress.com/comments-on-rwer-issue-no-96/>

One thing that most economists seem to agree on, regardless of their ideological persuasion or varying theoretical frameworks, is the importance of productivity increases. This “holy grail” is embraced by mainstream economists, heterodox, pluralist and dissident economists, policy makers and the general public as an obviously desirable goal that all economies must strive for. Yet of all the economic concepts widely in use, that of aggregate productivity in an economy may be the most problematic and full of conceptual and measurement holes.

This article is a critique of both the concept of productivity – especially labour productivity – as generally used in economic analysis, as well as the attempts to measure it to compare across countries and within a single country over time. It is essentially exploratory in nature: I will raise more questions than I can answer, and offer more criticism than solutions. Yet I hope to indicate that the widespread use of this concept is analytically and empirically flawed, and that a post-neoliberal economics that captures the true spirit of political economy must move beyond that to better, more relevant and “true” measures of human progress, even when progress is measured only in material terms.

At one level, the privileging of productivity appears to be so obvious that it requires no further elaboration. After all, productivity simply represents the amount of output per unit of input, and obviously it would be more “efficient”, less costly, and therefore presumably more desirable, to produce more output with the same or less inputs. But the first problem emerges immediately, in the very choice of variables. GDP, or national value added measured through national accounts data, is usually taken as the numerator. For international comparisons, there is the further choice of exchange rates for comparison, that is, whether to use Market Exchange Rates (those actually prevailing in any period) or Purchasing Power Parity (PPP) exchange rates. There are several concerns about GDP measures in themselves, and the extent to which estimated incomes represent even economic conditions.

The problems of using national income expressed in Gross Domestic Product are now widely recognised, in terms of the blindness to distributional issues and the inability to measure either the quality of life or the sustainability of any particular system of production, distribution and consumption. Despite these obvious limitations, however, it remains the most widely used indicator on any economy, and is generally the one that is tracked to determine both perceptions of national economic performance and policy orientations of most governments. This is unfortunate, because this obsession with GDP in itself, and independent of other markers of well-being, makes for problematic assessments of the actual performance of economies and, even more tellingly, for poor policy decisions and outcomes. Because GDP in most countries captures only marketed transactions, it excludes a significant amount of production of goods and services for self or household consumption. It makes market pricing the chief determinant of value, irrespective of the social value of any activity, which leads to massive undervaluation of what are now (especially post-pandemic) recognised as essential

¹ This paper is exploratory in nature, without providing very clear-cut answers and solutions to the problems identified. Comments, suggestions and further ideas will be greatly welcome.

social services relating to the care economy. It correspondingly overvalues those activities, goods and services that are priced higher because of the oligopolistic structure of markets. Because it does not estimate ecological and environmental costs in the inputs, it overestimates value added.

For example, a chaotic, polluting and unpleasant system of privatised urban transport involving a multiplicity of private and polluting vehicles on over-congested roads (as is common in many developing countries) typically generates more GDP than a safe, efficient and affordable system of public transport with lower vehicular congestion and a more pleasant living and working environment. In turn, where health services are commercialised, the consequent increase in morbidity from pollution and mortality from vehicular accidents also raises GDP, because of the resulting (largely private) expenditure on health services, etc. The deprivations caused by climate change and other evidence of ecological damage are the result of unsustainable patterns of economic activity that are simply not factored into estimated of national income, despite various attempts to incorporate them.

Services GDP is particularly hard to evaluate, because of the wrong valuation (from a human and social welfare standpoint) of different types of services. Financial services, for example, are hugely overvalued and over-rewarded, at least partly because of the political and lobbying power of financial interests in contemporary societies – and financial asset booms that reflect asset price changes then get reflected in increasing shares of financial services in national income, without any underlying real economic changes. Meanwhile, as the Covid-19 pandemic has shown, care services that are crucially important for human welfare, for the survival of societies and the resilience of economies, are routinely undervalued, with much of this activity performed unpaid (largely by women) within households or in extremely underpaid form. So GDP expansion as the desirable goal or indicator contains all sorts of concerns and contradictions.

Then there is the question of whether, even if GDP is accepted as the appropriate numerator, it is better to estimate “total factor productivity”, which is supposed to take account of all input use, or per worker productivity. While the former makes more logical sense, it is replete with concerns. Incorporating the role of other inputs like land, capital and intermediate inputs would make sense, but the valuation problems in each of these are immense, especially for assets like land and capital (and there is of course the inherent [contradiction identified by Piero Sraffa](#) of the self-referential nature of the measurement of capital and the rate of profit). All “total factor productivity” calculations are therefore suspect. Nevertheless, the “Solow residuals” emerging from such decomposition exercises, which are thereby supposed to represent the productivity improvement, are widely discussed and analysed, and have been variously ascribed to “social infrastructure” like institutions and government policies ([Hall and Jones, 1999](#)) and human capital and incentives for investment ([Acemoglu, 2001](#)).

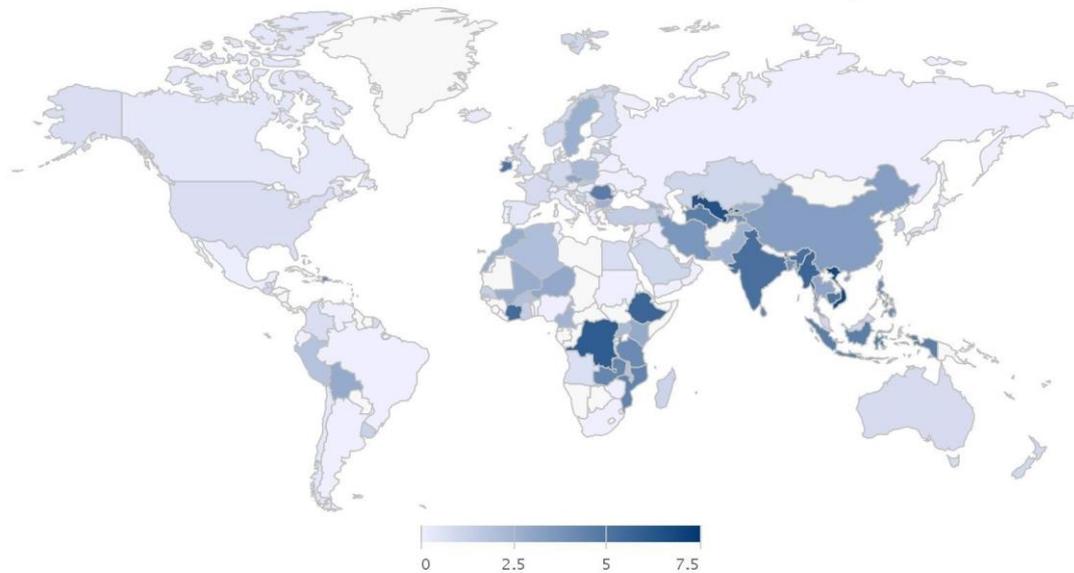
Because of problems in estimating total factor productivity, it has been more common to indulge in productivity comparisons across countries on the basis of per worker productivity or output per unit of labour, calculated in terms of hours worked. This is seen to indicate many other features of an economy: the per capita income; the extent of capital in use; the level of skill of the workforce; the potential of the economy to provide for the basic needs of the population; and so on. Consider the [ILO's definition of labour productivity](#): “Labour productivity measures the efficiency of a country with which inputs are used in an economy to produce goods and services and it offers a measure of economic growth, competitiveness, and living standards within a country.” Similarly, [the World Bank claims](#) that “labour productivity is used

to assess a country's economic ability to create and sustain decent employment opportunities with fair and equitable remuneration.”

Once again, there is no shortage of economists offering explanations for variations in output per worker across countries, as well as of changes in productivity over time within any country. Figure 1, for example, produced by the World Economic Forum, seeks to relate per capita incomes with changes in per worker productivity across the world. It suggests that lower-income countries display more rapid labour productivity growth, in accordance with a catching-up or convergence hypothesis.

Figure 1

Growth of GDP, Per Capita Income and Labor Productivity, 2015



*Note: Color ramp is based on Labor Productivity growth rates in 2015.
Source: The Conference Board Total Economy Database, May 2016.*

Source: <https://www.weforum.org/agenda/2016/07/what-is-productivity-and-how-do-you-measure-it/>, accessed on 15 July 2021

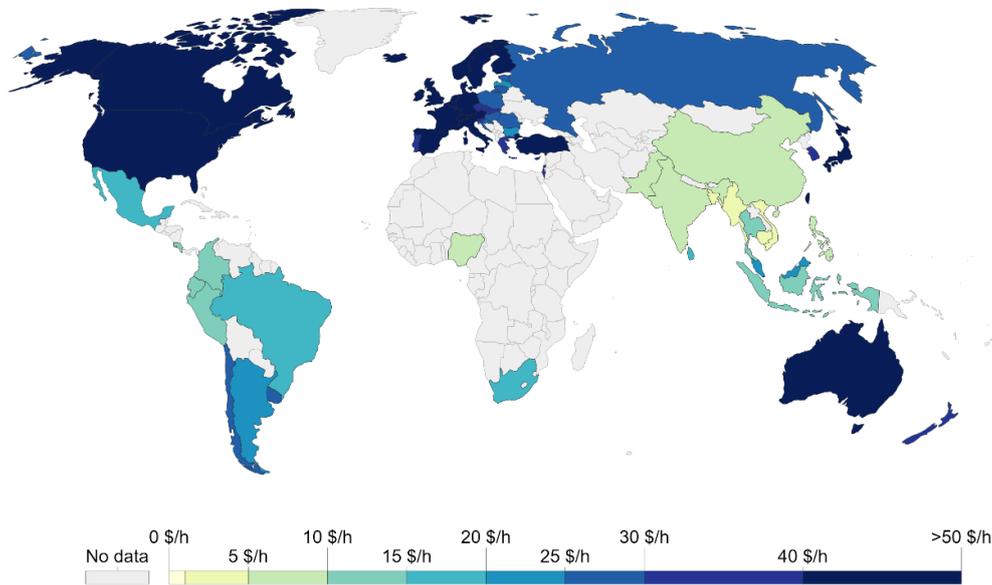
Obviously, per capita income in a country would be positively related with levels of productivity per worker (though not the rate of change) unless the worker population ratio changes over the period. But somehow this indicator of GDP per workers is taken to suggest something beyond this: the extent of technological advance of a country, its overall macroeconomic “efficiency”. Is this valid?

As already noted, the problem arises at two levels: in terms of the numerator and the denominator. With regard to the numerator, GDP, or aggregate value added in the economy, some concerns have already been expressed. The denominator, labour input, can be defined in terms of total hours worked or number of workers. Figure 2 provides some estimates of global differences of productivity per hour worked. It is immediately evident that data simply do not exist for a large part of the world. Even for countries where such estimates can be made, there are significant concerns about how accurate the estimates of hours worked are, especially in economies with a high degree of informality of the workforce and the presence of significant numbers of self-employed workers, where it is hard to gauge the actual hours of work.

Figure 2

Productivity per hour worked

Labor productivity per hour is measured as gross domestic product (GDP) per hour of work. GDP is measured in constant 2011 international-\$, which means it is adjusted for price differences between countries (PPP adjustment) and for inflation to allow comparisons between countries and over time.



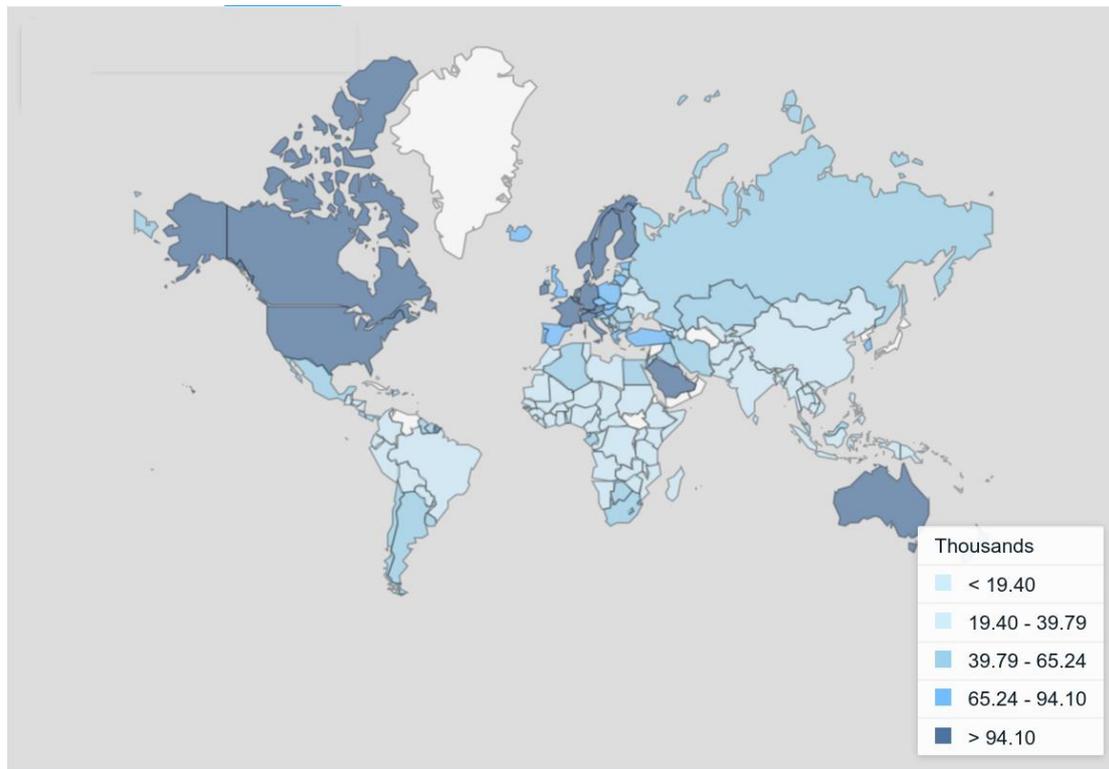
Source: based on Feenstra et al. (2015) Penn World Tables 9.1

[OurWorldInData.org/economic-growth](https://ourworldindata.org/economic-growth) • CC BY

Source: <https://ourworldindata.org/search?q=productivity> accessed on 13 July 2021

That is why productivity per worker has emerged as the most popular basis for cross-country productivity comparisons, as well as for assessments of changes in technology use and spread within economies over time. Figure 3 presents the World Bank's estimates of per worker productivity (GDP per person employed) based on national income calculated in terms of PPP exchange rates for 2017. Once again this appears to broadly track per capita GDP measures (in PPP exchange rates) which is not so surprising. But there are some anomalies and some outstanding issues, which become even more apparent when changes over time are tracked.

Figure 3



Source: <https://data.worldbank.org/indicator/SL.GDP.PCAP.EM.KD?view=map>, accessed on 15 July 2021

Table 1, which provides data on an arbitrarily chosen set of high, middle and low income countries, highlights some of the obvious anomalies, which in turn point to concerns with the very conceptualisation of this measure. Consider, to begin with, the absolute levels of per worker productivity in 2019 (taken as the latest year before the Covid-19 pandemic would have messed with the estimates). According to this set of estimates, in that year, economies like Armenia, Kazakhstan, Mexico and South Africa all had absolute productivity levels significantly higher than China. Other countries like Sri Lanka, Thailand, Colombia and Georgia also showed higher per worker productivity, which also flies in the face of evidence on external competitiveness, and other similar indicators. Meanwhile, Saudi Arabia shows extremely high per worker productivity – the second highest (after the US) in this list, and well above Australia, Canada, France or even Germany. The level of per worker productivity in Bahrain appears to be not so much lower than that in the United Kingdom.

Table 1. GDP per person employed (in constant 2017 PPP dollars)

	2000	2019	% change 2000-2019
Argentina	49,516	52,695	6.4
Armenia	11,575	39,786	246.4
Australia	81,602	99,569	20.1
Azerbaijan	8,673	30,271	243.7
Bahrain	1,08,537	71,898	-29.5
Bolivia	12,652	18,061	42.8
Cambodia	3,254	7,671	138.9
Canada	75,665	94,099	24.4
Chile	42,664	56,874	25.7
China	6,134	31,416	390.3
Colombia	24,882	32,610	24.1
France	98,524	1,03,185	12.5
Germany	95,083	1,02,107	10.3
Ethiopia	1,754	5,145	173.3
United Kingdom	81,474	84,206	14.2
Georgia	10,977	31,441	203.5
Ghana	7,094	13,349	86.8
India	7,285	19,270	168.7
Kazakhstan	22,769	57,620	148.4
Korea, Rep.	49,481	80,438	62.5
Malaysia	39,899	56,361	48.9
Mexico	46,269	44,839	-2.7
Mozambique	1,371	2,987	122.3
Myanmar	2,253	10,968	406.7
Nigeria	9,136	17,898	95.3
Saudi Arabia	1,42,881	1,16,313	-16.6
South Africa	39,296	43,893	11.5
Sri Lanka	15,467	33,767	121.1
Tajikistan	6,006	16,192	155.8
Tanzania	2,935	5,569	91.6
Thailand	18,065	31,680	85.3
United Kingdom	81,474	84,206	14.2
United States	1,00,390	1,31,047	27.1

Source: <https://data.worldbank.org/indicator/SL.GDP.PCAP.EM.KD?end=2020&start=1991&view=chart>

Rates of change are equally startling in some cases. The rapid increase in per worker productivity in China since the turn of the century may come as no surprise, given what we know about China's rapid rise and technological progress. But it also more than doubled in Georgia, went up by nearly two and a half times in Azerbaijan and more than four times in Myanmar! Meanwhile, per worker productivity showed an absolute decline of nearly 30 per cent in Bahrain over this same period.

Clearly, these figures are significantly affected by the numerator – GDP – which can change sharply for countries that depend excessively on certain mineral or raw material exports such as oil, as global prices change. This obviously plays a role in creating the dramatic increases and some declines in particular countries. These cannot be construed as saying anything about productivity as generally understood. Changes in per worker productivity in Saudi Arabia, Bahrain and even Myanmar could be at least partly understood in that light.

Another reason for these peculiar results is the use of PPP exchange rates to form the basis of comparison across countries. As I have written elsewhere ([Ghosh, 2018](#)) while PPP exchange rates appear to control for differences in price levels and standards of living in different countries, they are riddled with conceptual, methodological and empirical problems. They assume that the structure of each country's economy is similar to that of the benchmark country (the US) and changes in the same way over time beyond the reference year, which is clearly wrong across advanced and developing economies. The absence of weights within basic headings of goods and services, including the lack of representative weights, can result in these basic headings being priced using high-priced unrepresentative goods that are rarely consumed in some countries (Angus Deaton has provided the example of packaged corn flakes, which are available in poor countries, but only accessed by a relatively small minority of rich people). Country PPP rates are constructed from the prices of basic headings using expenditure weights from the national accounts – but these do not reflect the consumption patterns of people who are poor by global standards. While the current PPP measure does try to differentiate across regions, the different regions are linked using the region-wide 'super' PPP rates, which generate, for example, a price level for all of (say) Asia relative to the OECD countries—far too aggregative in a very disparate region to be at all accurate. There are additional concerns about the nature and coverage of the surveys that are conducted to establish the price levels in each country.

There is a further, and possibly even more damning, conceptual issue. In general, countries that have high PPP (that is where the actual purchasing power of the currency is deemed to be much higher than the nominal value) are typically low-income countries with low average wages. This occurs precisely because there is a significant section of the workforce that receives very low remuneration, which then means that goods and services are available more cheaply than in countries where the majority of workers receive higher wages. When even these activities are further subsidised by the widespread incidence of unpaid labour, as is typically the case in poor households in low income countries, then it is clear that the greater purchasing power of that currency reflects conditions of indigence and low or no remuneration for what could even be the majority of workers. Therefore, using PPP-modified GDP data may actually miss the point, by seeing as an 'advantage' (of greater purchasing power of a given monetary income) the very feature that reflects the greater absolute poverty of the majority of workers in an economy. This means that PPP income estimates effectively overstate incomes of poorer countries when it comes to comparing incomes across rich and poor countries. As countries move up the per capita income ladder, the difference between PPP and MER would reduce – not necessarily over time, but with increasing incomes of the

lower income country. As aggregate incomes increase, wages and prices in that economy also increase, typically relatively faster than in richer countries, thereby reducing the so-called 'PPP advantage'. This is strongly evident in the case of China, for example, where the ratio of per capita income measured in PPP terms to that measured in MER declined from 3.1 in 2000 to 1.7 in 2015 as the Chinese economy became richer.

All these factors inevitably make the use of PPP exchange rates in inter-country income comparisons extremely problematic. Certainly, they would not provide accurate estimates in comparisons of per worker productivity across countries.

But the denominator – the number of “workers” – is also a contentious issue. This is because it excludes the entire range of unpaid work that underwrites and typically subsidises the “paid” economy. Such work is largely (but not only) in subsistence provision and care activities within households and communities, and performed largely (but not only) by women and girls. The 19th International Conference of Labour Statisticians (ILO 2013) finally recognised this, by distinguishing between “work” and “employment” and expanding the concept of work: ‘Work comprises any activity performed by persons of any sex and age to produce goods or to provide services for use by others or for own use’. Employment – defined as ‘work for pay or profit’ – is therefore a subset of work.

This lack of recognition of a significant part of the work (dominantly provided by women) has several important economic and social implications. The unpaid-paid continuum of work serves to devalue both those who do it and the work they do. Thus for example, when women do enter labour markets, their wages tend to be lower than those of men – not only because they are willing to work for lower wages but because so much of their work is available for free. Related to this, the occupations in which women dominate tend to be lower paid – and the wage penalty extends even to men doing similar work, such as in the low paid care sector. Third, all this unpaid work provides a huge subsidy to the recognised economy and to the “formal sector”, which rely both directly and indirectly on the goods and services produced by these unsung and unrewarded workers. Because this contribution is not recognised, it could lead to measures suggesting rising aggregate labour productivity in the economy, which may be quite misplaced.

In other words, current estimates of per workers productivity across countries are poor indicators of the reality, because the conceptualisation and empirical estimation of both numerator and denominator are riddled with problems. How do we solve this? It's not clear what can be done within existing national accounts and statistical systems to make these problems go away: maybe the point is to search for a more reliable and valid indicator of human progress.

Author contact: [jayatijnu@gmail.com](mailto:jayati@jayatijnu@gmail.com)

SUGGESTED CITATION:

Ghosh, Jayati (2021) “Interrogating the holy grail of productivity growth.” *real-world economics review*, issue no. 96, 22 July, pp. 119-126, <http://www.paecon.net/PAERReview/issue96/Ghosh96.pdf>

You may post and read comments on this paper at <https://rwer.wordpress.com/comments-on-rwer-issue-no-96/>