How have global population and GDP per capita changed over human history?

Author: Toby Newberry

Introduction

This report surveys existing research on global population and global economic growth over the course of human history. These features are important from a longtermist perspective because they may inform forecasts of the likely future trajectory of human civilisation, and help us understand the relative effectiveness of different levers on civilisational progress. We can also use them to form a more accurate picture of how human civilisation has developed so far. They help us understand how many lives transpired in any given historical period, and the conditions in which they were lived.

My primary aim in writing this report was to identify the most accurate and extensive population and growth datasets available. I also had occasion to address a number of related questions. On the subject of global population, these questions included:

- Whether history of global population growth had just two important shifts, one at the agricultural revolution and one at the industrial revolution;
- Whether the adoption of agriculture increased the population growth rate;
- What rates of population growth were typical for hunter-gatherers, early agriculturalists, and populations after the industrial revolution; and
- How many people in total lived before the agricultural revolution.

On the subject of global economic growth, the most important additional question I considered was whether the history of economic growth can accurately be characterised as having only a single important transition, during the industrial revolution.

2. Methods & Data

I first conducted a broad search for potentially relevant sources. I looked for both datasets of global population or economic growth, and commentary on the relative merits of these datasets. I then read through the commentary to determine which dataset experts consider the most reliable. Finally, I summarised and wrote up the results.

For global population, I considered six datasets.[1] These were Clark (1967), McEvedy & Jones (1978), Biraben (1980), the UN (1999), HYDE 3.2 (2017), and Kaneda & Haub (2021). This section briefly describes each dataset, and makes some comments on other candidates that I didn't consider as closely. The next section discusses the relative merits of these datasets in greater detail.

Clark (1967) is a book-length discussion of global demographics and land use, which includes a range of estimates of the world population from 14 CE through to 1975. The approach taken is what I would describe as 'educated guess-work', which is broadly the same approach taken in all six candidate datasets. Helpfully, Durand (1977) includes a table containing Clark's estimates, and is easier to access than Clark's book.

McEvedy & Jones (1978) is another book-length treatment of global population questions. Its focus is almost exclusively on population estimates, and the book mostly consists of estimates of the populations of different parts of the world with graphs. The time period covered is BC 400 to AD 2000, and the method is also educated guess-work.

The remaining four datasets are more succinct. Biraben (1980) is a short paper that includes estimates of global population from BC 400 to AD 1970. UN (1999) is a report containing the UN's internal estimates of global population, which cover the period from BC 1 to projections for AD 2200. HYDE 3.2 (2017) refers to a database of global historical information, including population estimates from BC 10,00 to AD 2015. Finally, Kaneda & Haub (2021) is a set of global population estimates ranging from BC 50,000 to AD 2050, and produced by The Population Reference Bureau, an independent research centre for demographic questions.

In addition to these six candidate datasets, there are several other sets of estimates that are sometimes cited as authoritative, but that I didn't look into as closely. These include Durand (1977), Livi-Bacci (1992), Kremer (1993), and Maddison (2001). I didn't look into these because they are all essentially concatenations of the six datasets previously mentioned. Durand (1977) gives a table summarising estimates of global population, including Clark (1967) among other early attempts, but does not make novel estimates. Livi-Bacci (1992) uses the figures from Biraben (1980) for earlier dates, and UN figures for more recent dates (Livi-Bacci, p25)). Kremer (1993) primarily uses estimates from McEvedy & Jones (1978), supplemented with UN figures from 1920 onwards, and several estimates from Deevey (1960) for populations pre-10,000 BCE (Kremer, p 683). Maddison (2001) does contain some independent estimates, but uses McEvedy & Jones (1978) to 'fill holes' in the data, and draws heavily on their work for all pre-1500 estimates (Maddison (2001), p 230).

For questions of global economic growth, I considered two candidate datasets: De Long (1998) and the Maddison Project (ongoing). As far as I can tell, these are the only attempts to estimate global economic growth over the long run using something close to best practice techniques. De Long (1998) estimates global economic growth from BC 1 million to AD 2000. His method involves combining long run population estimates from Kremer (1993) with

[1] Our World in Data provided a useful stepping off point for this part of the inquiry.

Angus Maddison's estimates of GDP per capita for the 19th and 20th centuries. Essentially, De Long correlates these two sets of figures, then uses the correlation to retrodict GDP per capita for early history. The Maddison Project is an ongoing research initiative based at the University of Groningen, with the aim of continuing and updating Angus Maddison's estimates of economic performance indicators. It includes estimates of GDP per capita beginning in AD 1, divided into geographic regions.

3. Analysis & Discussion

3.1 Global Population

There are essentially three ways to estimate the historical population: using genetic evidence, using archaeological evidence, or making educated guesses.[2] In practice, however, the first two methods do not yield results that are sufficiently reliable, or sufficiently comprehensive, to be useful in the context of this inquiry.

Genetic evidence can be used to estimate historical populations as follows. First, you estimate the dates of origin for a variety of mutations in the current gene pool, by looking at factors like the geographic spread of those mutations. You then estimate the relative contributions that different time periods made to the current gene pool. Together with some additional assumptions, these estimates can then be converted to estimates of population. To give an illustrative example, if you found that many present-day mutations date to around BC 20,000, while only a small number date to around BC 15,000, this could imply that the global population declined during that period. However, as noted, this method is not yet sophisticated enough to give results that are suitable for addressing the large scale questions considered here.

The archaeological approach is comparatively straightforward, and hinges on the idea that larger populations will tend to leave larger archaeological footprints. Simply put, if there were many more people living at time A rather than time B, we would expect to find a greater density of bones, tools, pots, or other archaeological residue dating from time A. Unfortunately, this approach faces an array of methodological challenges, including the fact that some sorts of deposits tend to disappear over time, and that the size of an archaeological footprint is sensitive to many other factors as well as population. Some of these challenges are surveyed in Williams (2012). More generally, the archaeological record can support some population estimates about specific regions at specific times, but is too patchy to support estimates of global population over many millennia.

M estimates in the literature use educated guess-work. At its best, this approach involves estimating global population from many regional estimates, drawing on domain experts or published figures for the regional estimates, and making principled adjustments in light of qualitative arguments. At its worst, this method involves fudging figures to match proposed historical narratives, making non-robust guesses, or making highly questionable inferential

[2] My discussion of these methods draws heavily on notes made on these subjects by Ben Garfinkel and David Roodman.

leaps. As an example of the latter, McEvedy & Jones (1978) estimate the global population before BC 10,000 based partly on the density of present-day gorilla populations (p13-14).

More commonly, a central problem with this approach is that it usually involves making small modifications to previous estimates, which are themselves modified versions of even earlier estimates. Because of this, many relatively recent estimates actually incorporate figures from 40-100 years ago, and very few, if any, of the sets of estimates can properly be considered independent. As a more acute form of the problem, many authors use at least some of the figures from McEvedy & Jones (1978) (e.g. Kremer (1993), Maddison (2001), HYDE 3.2 (2017)), while McEvedy & Jones themselves lean heavily on estimates made decades earlier (especially Russell (1948)).

A further problem with the guess-work approach is that questions of global population were somewhat politicised throughout the 20th century. For example, it is possible that population estimates were influenced by concerns about overpopulation or the idea that certain parts of the world were ‘terra nullius’ before European conquest.

In light of these problems, one conclusion of this report is that the quality of all long-run estimates of the global population is surprisingly poor. While there are some topics about which we can make robust claims over long timescales, such as claims about Earth’s climate based on evidence from ice-cores, global population is not yet among these topics. This should serve as a caveat that the specific claims made below come with a high degree of uncertainty. I should also note that many of the authors mentioned so far are fully aware of the speculative nature of their work. McEvedy & Jones even joke that on rare occasions they just “pulled figures out of the sky,” (McEvedy & Jones, 1978, p10).

Of the six candidate datasets given in the previous section, I think HYDE 3.2 (2017) provides the most accurate estimates of global population over the period 5,000 BCE to the present day. The reasons for this verdict are as follows.

Both Clark (1967) and Biraben (1980) suffer from more severe qualitative problems than do the other datasets. Clark (1967) has been the subject of several reasonable critiques. For instance, Durand (1977) grades all of Clark’s population estimates on an A to D scale, and assigns most of them D grades, meaning they are based on assumed rates of population growth, rather than any specific data relating to the time period or region in question. With the exception of estimates for the population of China, none of Clark’s estimates from before 1750 receives a grade higher than C (see Durand (1977), p255 for an explanation of his grading system). Maddison (2001) also compares Clark and Biraben’s estimates unfavourably to McEvedy & Jones (1978), noting that they are less detailed and well-documented, and involve judgements with which he disagrees. Biraben (1980), for his part, is relatively clear in stating that his estimates should not be considered properly authoritative:

> “We had neither the means nor the time to go over all the existing documentation in order to obtain new complementary data. We therefore limited ourselves to collating the published figures for each country; then, taking into consideration what is known about the economic, political, and medical history of each country, we drew a curve through these figures and extended it into the past.” (Biraben, 1980, p655-656).

Both Kaneda & Haub (2021) and the UN (1999) suffer from issues of transparency; in neither case was I able adequately to account for the sources of the relevant figures. The UN (1999) includes a table of population estimates going back to AD 1, and cites several other pieces of work by the UN. However, in consulting these other sources I was unable to find population estimates going back further than AD 1950, and the origin of the earlier figures remains mysterious. Kaneda & Haub (2021) offer some insight into their methodology, which was to produce population estimates for certain benchmark years, and then combine these with birth rate estimates to generate a continuous series. However, as the source for their historical population estimates they cite Dudley Poston Jr, a demographer from Texas A&M University, and do not offer details on how Poston Jr. came to his estimates.[3]

McEvedy & Jones (1978) and HYDE 3.2 (2017) are very similar datasets. In fact, the HYDE dataset is essentially identical to McEvedy & Jones, except that it incorporates more robust data from the UN for the mid 20th century onwards, and supplements some earlier figures with independent subnational estimates. This last comment is slightly opaque, in that I was unable to work out exactly which estimates it refers to. That said, my view is that on balance the HYDE data is slightly preferable to McEvedy & Jones, if only because it includes more reliable recent figures. Both HYDE and McEvedy & Jones are preferable to the other four candidate datasets, for the reasons given above.

One question that can be asked about global population concerns transitions between different growth modes. Specifically, it is often suggested that the history of global population growth contains two important shifts, one occurring at the dawn of agriculture, and one at the industrial revolution.

The HYDE 3.2 data broadly supports this claim, but there are a number of clarifications worth making. First, the claim that the history of global population contains two transition points, or two shifts between different growth modes, could be interpreted as implying a piecewise structure to the data. This is not the interpretation I am taking. Instead, I understand it, the claim is that the history of global population looks more or less like it has three parts, with transitions between those parts occurring around the time of the dawn of agriculture, and around the time of the industrial revolution. This is what the HYDE 3.2 data shows. In particular, it shows that global population was approximately level before around BC 10,000, rising quite slowly between BC 10,000 and around AD 1700, and rising sharply from that point on. Note, however, that this does not necessarily imply that ‘something special’ happened during the transition times.

The first transition, occurring at the dawn of agriculture, is also well-supported by evidence from other sources. In archaeology, for example, this period is known as the ‘neolithic demographic transition’, and there is evidence for a sharp increase in birth rate[4] following the advent of agriculture (see, e.g. Bocquet-Appel & Bar-Yosef, (2008), Bocquet-Appel (2002),

[3] They cite the UN and previous work by the Population Reference Bureau for more recent figures (from about 1950 onwards).

[4] Note that an increase in birth rate does not necessarily imply an increase in population, since it could be offset by an increase in mortality. I take it that evidence for the former without additional evidence for the latter should still be considered evidence in favour of a population increase.

Bocquet-Appel & Paz de Miguel Ibanez (2002)). There is also direct genetic evidence for a sharp increase in population at this time. For example, Gignoux, Henn, & Mountain (2011) find that “the invention of agriculture facilitated a fivefold increase in population growth relative to more ancient expansions of hunter-gatherers.”

The second transition, occurring at the industrial revolution, is less well-supported by additional evidence, but still appears plausible. The second graph shown in Section 4, which uses the HYDE 3.2 data, displays a sharp uptick in global population around AD 1750. In addition, at least one prominent economist has publicly endorsed the view that there was an “unprecedented increase in population growth during the early stages of industrialisation,” (Galor, (2005)).

There are several other questions about global population that were mentioned in the introduction, and which can be answered using the HYDE data together with other resources. To begin with, the HYDE data indicates that the population growth rate was higher after the advent of agriculture than in the hunter-gatherer era. More specifically, the data shows essentially no sustained growth in population for hunter-gatherers, with typical annual growth of around 0.1% for agriculturalists.[5] One caveat to this is that pre-agricultural populations almost certainly rose and fell; that is, the lack of growth should not be taken to mean that the pre-agricultural population was constant. After the industrial revolution, the population growth rate rose again to between 0 and 2% per year, with a typical rate of around 1%.

HYDE 3.2 also includes estimates of typical population densities for agriculturalists and people living after the industrial revolution, which are a helpful complement to the estimates of growth rate. These figures are summarised in the table below, though it should be noted that the density estimates for hunter-gatherers come from McEvedy & Jones (1978), and should be considered highly uncertain.

	% growth (annual)	**Population density (people/km^-2)**
Hunter-gatherers	Typical: 0% Range: not enough reliable data	Typical: 0.1 Range: 0.0 to 3.5
Agriculturalists	Typical: 0.1% Range: -0.06 to 0.5%	Typical: 0.6 Range: 0.03 to 6.0
Post-industrial revolution	Typical: 1% Range: 0 to 2.2%	Typical: 20 Range: 6 to 56

We might also wonder about the total number of people who have *ever* lived. While the HYDE 3.2 data offers estimates of the total number who have lived since around 10,000 BCE, it is silent on the question of how many people lived before the agricultural revolution.

[5] Here I am interpreting ‘agriculturalists’ quite broadly, as anyone living between 10,000 BCE and 1750 CE. This is so that the three categories are jointly-exhaustive: ‘hunter-gatherers’ covers the period before 10,000 BCE, and ‘post-industrial revolution’ covers the period after 1750 CE.

Fortunately, this question is addressed by Kaneda & Haub (2021), who estimate that the total number of people who lived between BC 190,000 and BC 10,000 was between 8 and 9 billion. Note again that this figure is a personal estimate made by a sociology professor, and I am unsure of the precise methodology at work. Note also that the cutoff date of 190,000 BCE is somewhat arbitrary, since there is not a well-defined point at which pre-human ancestors became fully-fledged Homo sapiens.

3.2 Global GDP per capita

Of the two candidate datasets for global GDP per capita, my view is that the ex-Nordhaus series included in De Long (1998) is the most accurate. The reasons for this are as follows.

The Maddison Project data has several positive features. In particular, their estimates are about 25 years more recent than De Long's, and appear to represent considerably more person-hours of research; they have been updated relatively regularly since 1995. However, there are two features of the dataset that together make it unsuitable for our purposes. First, the Maddison Project estimates do not extend any earlier than AD 1. This would not be especially problematic on its own, since we could simply use De Long's data for earlier years. Unfortunately, the Maddison Project data is also separated into geographic subregions, and there is no clear and robust way to convert its figures into an estimate of world GDP. One might think these subregions could simply be summed up, but the result of doing so has some unusual features[6] that imply the geographic regions are not comprehensive (that is, there might be 'interstitial regions' that do not have estimates included in the original dataset, and which change over time).

De Long (1998) in fact consists of three separate datasets. One of these is based on Maddison's modelling assumptions[7], one on De Long's own model, and one on De Long's model with an additional important assumption (specified below). Notably, Maddison's model assumes that "pre-1500 GDP per capita was constant at near-Malthusian bio-cultural subsistence." De Long describes his own model as a "fitted relationship between population growth and Maddison-concept GDP per capita." As noted in Section 2, he uses Maddison's GDP estimates from 1820 onwards to establish a correlation between population and GDP per capita, and then extends this backwards in time using population estimates from Kremer (1993). The additional assumption included in the third dataset is that GDP per capita should be multiplied by a factor of four for the period from 1800 onwards. This is to account for the fact that the range of goods and services on which one can spend money has increased over time - though De Long admits the precise implementation is somewhat arbitrary.

In my view, the De Long model without the additional assumption, which he calls 'the ex-Nordhaus series', is the most accurate of the three. The Maddison series can be excluded on the grounds that it assumes a particular answer to most of the question without

[6] For example, it shows a sharp uptick around 1250 CE, followed by a slow downward trend until around 1700. As noted, these appear to be an artefact of the method, rather than genuine features of the trend in GDP per capita over time.

[7] These are as follows: "...assume that GDP per capita was constant in Asia and Africa from 1500-1820, grew at 0.1 percent per year from 1500-1820 in Latin America and Eastern Europe, and grew at 0.2 percent per year from 1500-1820 in Western Europe… [and] that pre-1500 GDP per capita was constant at near-Malthusian bio-cultural subsistence."

argument. I have two reasons for preferring De Long's unmodified series to the series with the additional assumption. First, as stated, the four-fold multiplication is openly arbitrary. I therefore think it is more forthright to present the data without the extra assumption, and simply note that a possible limitation of the method is that it ignores 'Nordhaus effects'. Second, I do not find the explanation for these effects entirely convincing. De Long writes that he "would be extremely unhappy if [he] were handed [his] current income, told that [he] could spend it on goods at current prices, but that [he] was prohibited from buying anything that was not made before 1800." This strikes me as misleading, in that the relevant comparison should also include De Long having had his memories replaced with someone who grew up around 1800 (say). More generally, I think the intuition that the present cornucopia of goods and services effectively translates to greater GDP per capita assumes a low degree of hedonic adaptation. It is not obvious to me that today's relative abundance actually means people live at a higher quality of life.

As with global population, one interesting question that can be asked of the data on global economic growth concerns transition points. In particular, we might wonder whether the history of global economic growth contains only a single transition point, occurring at the industrial revolution. On the basis of both the De Long data and the Maddison Project estimates, I think this claim is justified. Graphs of GDP per capita over time tend to be mostly flat until around 1800, when they have a sharp uptick. This is true of the global data, as well as most country-specific data (see the Maddison Project (ongoing) and Broadberry (2013)). Importantly, it contrasts with observed trends in global population. In looking at the graphs of global population in the next section, and extending the first graph as a flat line to the left, it is clear that 'something changed around BC 10,000, and perhaps something changed again around 1800'. In looking at the graph of global economic growth, on the other hand, there is no equivalent transition around BC 10,000.

Another question one might have about global economic growth concerns what sort of model best fits with the estimated data-points. For example, it could be that the estimated and known data best fits with an exponential model, a hyperbolic model, a series of exponentials, or something else entirely. I will not attempt a comprehensive answer to this question here, but instead note that there is a reasonable and ongoing debate about it, and sketch some of the positions in that debate. In brief, Oded Galor's Unified Growth Theory holds that the history of economic growth consists of two phases, the first involving a long period of Malthusian stagnation, followed by a period of sustained economic growth beginning at the industrial revolution (see, e.g. Galor (2005)). Robin Hanson has argued for a view in which the growth rate increased during the agricultural and industrial revolutions, but was on average neither increasing nor decreasing for the intervening millennia. He also presents a model of long-term growth as a series of three exponentials, corresponding to the pre-agricultural, agricultural, and industrial eras (Hanson, (2000)). Kremer (1993) makes the claim that the economic growth rate has risen hyperbolically, or 'in proportion to the size of the economy'.[8] More recently, Ben Garfinkel conducted a useful investigation of Kremer's case (Garfinkel, (2020)). David Roodman has also done technical work on similar questions, incorporating stochasticity into a base-hyperbolic model (Roodman, 2020).

[8] Note that these theories refer to growth, or total output, rather than GDP per capita.

4. Conclusions

Based on the reasoning given above, my preferred graphs of global population use the HYDE 3.2 dataset, and look like this:

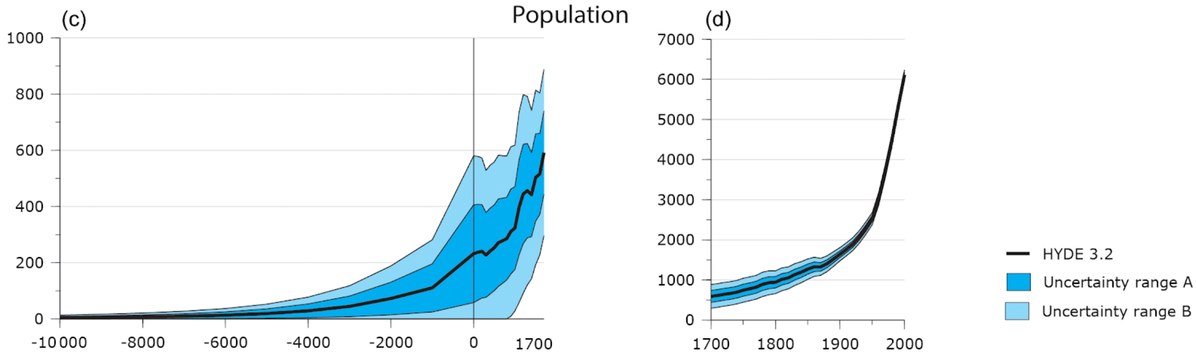

The first graph covers the pre-industrial period, and exhibits the neolithic transition described in Section 3.1. If we imagine extending the graph as a horizontal line leftwards,[9] then it looks as if the graph begins its gradual upward trajectory shortly after 10,000 BC. Together, the two graphs also capture the idea that our information about global population is highly uncertain until about 1900, from which year we have more reliable census data.[10]

The second graph slightly obscures the transition at the industrial revolution, since it uses a different x-axis. However, this second transition does show up as a sharp uptick on the far right of the graph when the full dataset is displayed, as below.

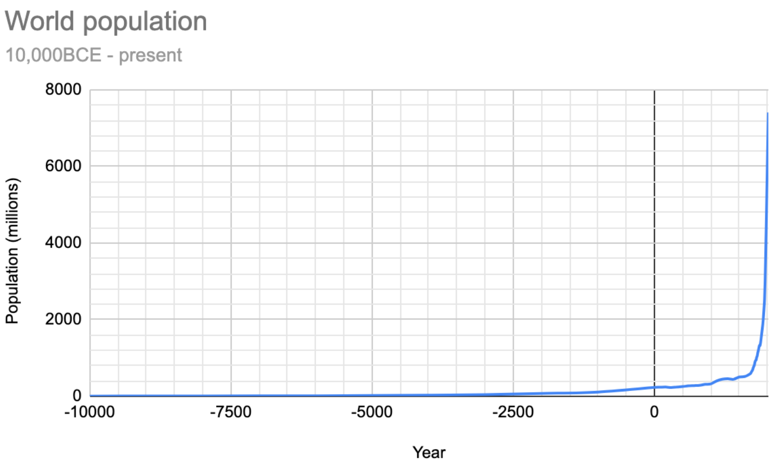

Using the ex-Nordhaus series from De Long (1998), I think the most accurate graph of world GDP per capita over time looks like this:

---

[9] Here I assume there were no extremely large populations of early humans of which we are unaware.

[10] Here is an explanation of the uncertainty ranges: "These uncertainty ranges were partly based on the ranges we could find in literature and partly on our own expert judgement and should be treated with care. The uncertainty range A is cautiously estimated at 1 % in 2000 CE, 5 % in 1900 CE, 10 % in 1800 CE, 25 % in 1700 CE, 50 % in 850 CE, 75 % in 1 CE and 95 % in 10 000 BCE.," (Goldewijk et al, (2017), p943). Uncertainty range B is simply twice range A.

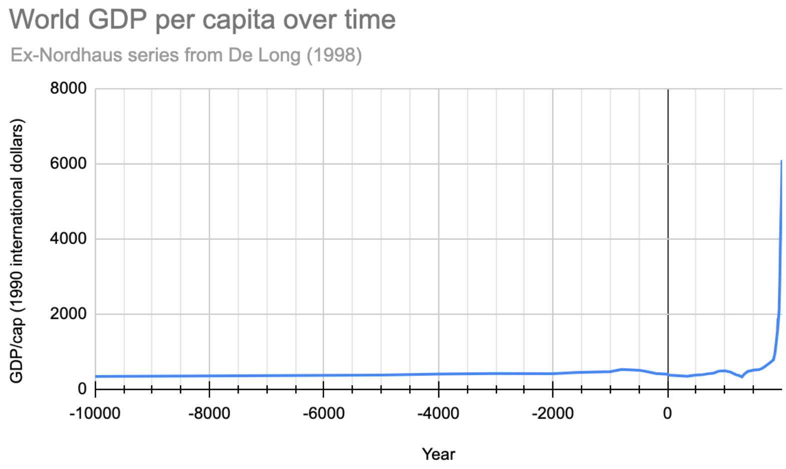

Each of these minor gridlines on the horizontal axis indicates 500 years, so that the sharp rise on the right hand side of the graph occurs at approximately 1750 CE.

Along with these graphs and their associated datasets, the foregoing discussion reached a number of other conclusions. In particular:

(1) the quality of all long-run estimates of global population is surprisingly poor;
(2) it is fair to characterise the history of world population growth as containing two important transitions, one occurring at the dawn of agriculture and one at the industrial revolution;
(3) It is fair to characterise the history of growth in GDP per capita as containing one important transition, occurring at the industrial revolution; and
(4) between 8 and 9 billion people lived before the agricultural revolution.

5. References

Biraben (1980): Biraben, J. N., ‘An Essay Concerning Mankind’s Demographic Evolution’, *Journal of Human Evolution*, Vol. 9, No. 8, 1980.

Bocquet-Appel (2002): Bocquet-Appel, J-P., ‘Paleoanthropological Traces of a Neolithic Demographic Transition’, *Current Anthropology*, Vol. 43, No. 4, 2002.

Bocquet-Appel & Bar-Yosef (2008): Bocquet-Appel & Bar-Yosef, *The Neolithic Demographic Transition and its Consequences*, Springer Science+Business Media, 2008.

Bocquet-Appel & de Miguel Ibáñez (2002): Bocquet-Appel, J-P., and de Miguel Ibáñez, M P., ‘Demografía de la difusión neolítica en Europa y lo datos paleoantropológicos’, *Saguntum*, Vol. 5, 2002.

Clark (1967): Clark, C., *Population Growth and Land Use*, The Macmillan Press Limited, London and Basingstoke, 1967.

Deevey (1960): Deevey, E. S., ‘The Human Population,’ *Scientific American*, CCIII, 1960.

De Long (1998): De Long, J. B., ‘Estimates of World GDP, One Million B.C. – Present’, Department of Economics, UC Berkeley, 1998.

Durand (1977): Durand, J., ‘Historical Estimates of World Population: An Evaluation’, *Population and Development Review*, Vol. 3, No. 3, 1977.

Galor (2005): Galor, O., ‘From stagnation to growth: Unified Growth Theory’, in P. Aghion & S. Durlauf (Eds.), *Handbook of Economic Growth*, Amsterdam: Elsevier, 2005.

Garfinkel (2020): Garfinkel, B., ‘Does Economic History Point Toward a Singularity?’, *Effective Altruism Forum*, 2020.

Gignoux, Henn, & Mountain (2011): Gignoux, Henn, & Mountain, ‘Rapid, global demographic expansions after the origins of agriculture’, *Proceedings of the National Academy of Sciences*, Vol. 108, No. 15, 2011.

HYDE (2017): Goldewijk et al, ‘Anthropogenic land use estimates for the Holocene – HYDE 3.2’, *Earth System Science Data*, No. 9, 2017.

Kaneda & Haub (2021): Kaneda & Haub, ‘How Many People Have Ever Lived on Earth?’, *Population Reference Bureau*, 2021.

Kremer (1993): Kremer, M., ‘Population Growth and Technological Change: One Million B.C. to 1990’, *Quarterly Journal of Economics*, August 1993.

Livi-Bacci (1992): Livi-Bacci, M., *A Concise History of World Population*, Blackwell Publishers Ltd, 1992.

Maddison (2001): Maddison, A., *The World Economy: A Millennial Perspective*, Paris: OECD, 2001.

The Maddison Project (ongoing): Maddison Project Database 2020, University of Groningen, accessed at https://www.rug.nl/ggdc/historicaldevelopment/maddison/releases/maddison-project-database-2020.

McEvedy & Jones (1978): McEvedy & Jones, *Atlas of World Population History*, Facts on File, New York, 1978.

Our World in Data: Roser, M., Ritchie, H., and Ortiz-Ospina, E., ‘World Population Growth’, accessed at https://ourworldindata.org/world-population-growth.

Roodman (2020): Roodman, D., ‘On the probability distribution of long-term changes in the growth rate of the global economy: An outside view,’ *Open Philanthropy*, 2020.

Russell (1948): Russell, J., *British Medieval Population*, The University of New Mexico Press, 1948.

The UN (1999): The UN, ‘The World at Six Billion’, Population Division, Department of Economics and Social Affairs, 1999.

Williams (2012): Williams, A., ‘The use of summed radiocarbon probability distributions in archaeology: a review of methods,’ *Journal of Archaeological Science*, Vol. 39, Issue 3, 2012.