Historical incomes on the cheap: a rough and ready shortcut to global GDP comparisons

Tom Westland[1]

[1]University of Cambridge

January 10, 2021

There are essentially two kinds of historical GDP reconstructions: those that are interested in growth, and those that are interested in cross-country comparisons. This paper is one of the latter. I make estimates of GDP per capita for about 100 countries c.1900, roughly doubling what is available in the current Maddison Project database. I use only three main sources of data: grain wages, urbanisation rates, and commodity trade. The first and third variables have been taken from a new dataset of global commodity trade and wages for the period of the First Globalization (Weber, Semeniuk, Westland and Liang, 2021). The second is taken from ClioInfra, though modified and added with secondary sources in some cases.

Benchmarking using the Malanima shortcut.

My method is an adaptation of the income benchmarking method used by to estimate GDP per capita for African countries for 1950.(Bolt et al., 2018) In this section, I briefly lay out their method, and discuss some potential problems in applying this methodology more broadly. In the following section, I outline solutions to these problems.

Bolt et al. essentially use a variation on what we might call the 'Malanima shortcut' (Malanima, 2011). Malanima's method departs from two ideas: one is that agricultural consumption per capita can be derived using a simple demand function for food:

$$Q_A \quad = f = W^{\alpha} \cdot P_A^{\beta} \cdot \; P_N^{\gamma}$$

where w is the real wage, P_A the price of agricultural goods, P_N the price of non-agricultural goods, and the Greek exponents represent the respective elasticities. Agricultural output (which is what we are interested in) might be supposed to be equal to consumption, as it is in Malanima's work; one might also suppose it to be some multiple of agricultural consumption: greater than 1 if the economy in question is a net agricultural exporter, and less than one if it is a net agricultural importer.

Once agricultural output has been obtained, the next major problem is to determine non-agricultural output per capita. Malanima offered two possibilites: once based on the (linear) relationship between urbanisation and the non-agricultural sector's share in output observed in periods for which there is independent data on these, and the other based on more scanty data on the relationship between the non-agricultural sector's share of employment and its share of output. These relationships are then applied to previous periods without adjustment.

Because the Maddison Project needed a cross-sectional estimate, they adapted Malanima's method, using a Törnqvist index to estimate the relative income of country *i* relative to a some reference point (which I will henceforth call an 'anchor country'—in the case of Bolt et al, South Africa was used:

$$\log\left(\frac{Q_i}{Q_j}\right) = \frac{1}{2}\left(\frac{P_i^a \cdot Q_i^a}{P_j^a \cdot Q_j^a}\right)\log\left(\frac{Q_i^a}{Q_j^a}\right) + \frac{1}{2}\left(\left(\frac{P_i^n \cdot Q_i^n}{P_j^n \cdot Q_j^n}\right)\log\left(\frac{Q_i^n}{Q_j^n}\right)\right)$$

where *Q* denotes the volume of output, *P* denotes prices, *a* and *n* denote agricultural and non-agricultural output respectively, *i* indexes countries and *j* denotes the anchor country. Due to the paucity of data on agricultural and non-agricultural production in developing countries, several shortcuts are used to estimate these quantities.

As above, agricultural productivity per person is derived using a demand function, though the relative prices of agricultural and non-agricultural goods are abstracted from, such that the only determinant of demand for food is the real wage:

$$Q_a = Nw^\beta \therefore q_a = w^\beta$$

To get from agricultural output to non-agricultural output requires strong assumptions. Whereas Malanima relies on a historical relationship between sectoral shares and urbanisation, Bolt et al. use urbanisation rates to estimate *employment* shares and then use a constant productivity gap *p,* to arrive at output in the non-agricultural sector:

$$Q_n = pQ_a\frac{u}{u-1} = pNw^\beta\frac{u}{1-u}$$

These estimates can then be plugged into equation 2 to estimate GDP relativities. If we know GDP per capita for the anchor country *j*, then we can use multiply it by the ratio to obtain GDP per capita in country *i*. In the followings section, I outline several modifications.

Modifications to the Maddison-Malanima method

There are two major objections one might raise to this procedure. The first concerns the productivity constant *p;* the second is to do with the way in which food production is estimated. I take each objection in turn.

The productivity gap

Bolt et al. use a constant productivity gap, $p = 2$, taken from a recent analysis of the agricultural productivity gap (Gollin et al., 2014). Gollin, Lakagos and Waugh (henceforth, GLW) estimate that once you account for differing levels of human capital and for different working hours, agricultural productivity is on average half that of non-agricultural productivity. One initial objection to the choice of $p = 2$ is simply that adjusting for human capital and working hours is not justified if p is to be used to compute output. In GLW, the purpose is to explain the sectoral productivity gap, some of which can be explained by accounting for the characteristics of workers in each sector. But our purposes here are not explanation but rather prediction—we are interested, in other words, in the size of the gap, not its causes. If we were to choose a constant gap, therefore, the 'raw' gap of 3 would seem more appropriate than the adjusted gap of 2.

This solution brings with it new problems, however. While the adjusted gaps are reasonably tightly distributed around the mean, 2, this is not the case for the unadjusted gaps, as we can see in Figure 1. Moreover, the use of a constant gap of $p=2$ for all countries is justified in Bolt et al. on the basis that GLW find that "the labor productivity ratio does not vary systematically with the country's income level." This seems more or less true for the adjusted gaps, used by the Maddison estimates; however, if we decide to use the raw gaps, there is a more clearly negative relationship between p and income, at least at lower income levels. However, it is impractical to use this negative relationship to establish appropriate choices of p for countries—to do so would be quite literally to beg the question. We would have to know a country's GDP in order to know its value of p, and GDP is the thing we are trying to estimate. In any case, while the negative relationship exists, it is not a particularly powerful predictor of the agricultural productivity gap.

Were it possible to identify what the relevant geographical endowments were, we could perhaps obtain more historically sensitive estimates of p. Unfortunately, the origin of the agricultural productivity gap is still mostly a mystery: while GLW identify several geographical variables that might explain the gaps, their results are relatively fragile, and the existence of multicollinearity makes it hard to distinguish which of the variables may have an ultimately causal role. Indeed, in their 'kitchen sink' regression, only ruggedness of terrain (Nunn and Puga, 2012)—is a statistically significant predictor of the agricultural productivity gap, and the R^2 obtained is low, at 0.18. One variable that GLW do not test, but which is highly correlated with the raw productivity gap is the share of labour in agriculture (the relationship is depicted graphically in Figure 2). This relationship is fairly unremarkable, theoretically: in societies that are heavily rural, one would expect that those people that leave agriculture are precisely those people that possess the kinds of skills that are likely to make them highly productive in non-agricultural work. In more developed societies, however, the marginal product of labour is more likely to be equalised across the sectors.

This relationship does not do anything to illuminate the origins of the agricul-

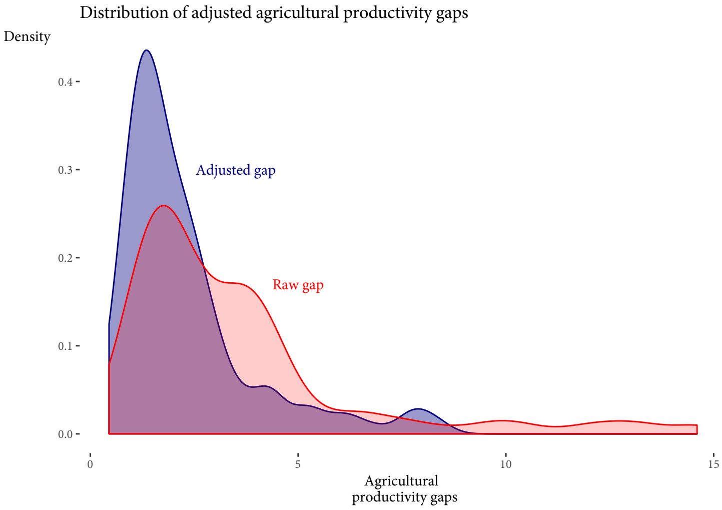

Figure 1: The distribution of agricultural productivity gaps. In blue are gaps adjusted for human capital of workers and hours worked. In red are the unadjusted gaps.

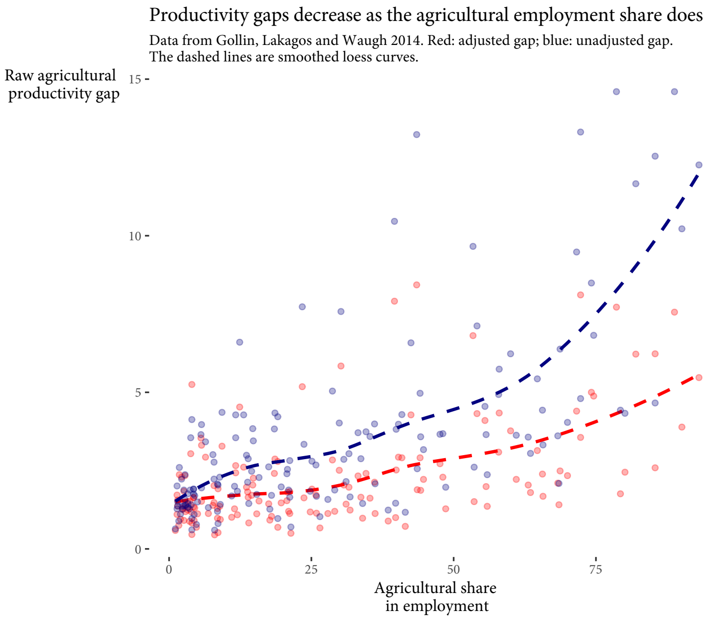

Figure 2: Relationship between agriculture's share of employment and the raw and adjusted agricultural productivity gaps

tural productivity gap, other than to underline that it is probably a result of the barriers to the sectoral reallocation of labour. This explains why GLW did not include it as a right-hand side variable in their regressions. However, if all we wish to do is estimate *p*, rather than explain it, the relationship is useful: and, since Bolt et al. have already defined the agricultural share of labour to be 1 – *u*, any historical work would not require any further data collection. A simple bivariate OLS regression of employment shares on the productivity gap using contemporary data yields quite strong results, with an R^2 of over 0.50. However, the relationship appears to be non-linear. In Table 1, I estimate the relationship adding a quadratic and a cubic term, then I add a variable that measures the share of fuels and metals in the country's export basket in 2005, taken from the World Bank database. Finally, in specification 3, I add an interaction term between the agricultural employment share and the fuel and metal share. The R^2, at just over 0.6, is lower than one might like for the purposes of prediction, but the regression is more informative than simply taking the mean value. Though in future we can hope for more historically grounded estimates of the productivity gap, in what follows, I take the coefficients from regression 3 and use them to estimates values of *p* in the historical period.

Mechanically, the effect of this procedure will be to increase the estimated GDP per capita of poor urbanised countries relative to poor non-urbanised ones. What is encouraging is that all of the variables in regression (3) are plausibly obtainable for most countries in the early twentieth century. For our sample of countries in the early twentieth century, I take the agricultural share of employment as equal to 1 minus the urbanisation rate, taking the urbanisation rate from Clio-Infra. I estimate the fuel and metal share of exports in the same way as the World Bank.

The balanced agricultural trade assumption

Implicitly, we assume that food supply and food demand are equal without any need for external trade. But this is unrealistic; some countries are net food importers; others are net food exporters. We must also deal with the fact that farmers do not merely produce staple foodstuffs (which is the implicit assumption in the Malanima method). They can also produce cash crops, as many did in the Belle Epoque. Therefore we need to take account of trade in agricultural goods. I deal with each of these probems separately.

First, suppose that there are two kinds of worker (rural and urban) and three kinds of good: a staple food product (call it maize), a non-staple agricultural product destined mainly for export (call it cocoa) and a non-agricultural good (call it machines). The demand for food is as above: i.e., it is a function of the real wage: $d_F = w^\beta$. Let us assume both the wage, and hence the level of food demand, are fixed. In an economy without food trade, rural

Table 1: Predictors of the agricultural productivity gap

	Dependent variable:		
	raw		
	(1)	(2)	(3)
Agriculture share	0.104*	0.107**	0.120**
	(0.053)	(0.051)	(0.050)
Squared	−0.002	−0.002	−0.003**
	(0.002)	(0.002)	(0.002)
Cubed	0.00003**	0.00003**	0.00004***
	(0.00001)	(0.00001)	(0.00001)
Fuel and metal exports		0.019***	0.002
		(0.006)	(0.009)
Interaction			0.001**
			(0.0003)
Constant	1.475***	1.108***	1.310***
	(0.418)	(0.419)	(0.420)
Observations	132	132	132
R^2	0.589	0.620	0.636
Adjusted R^2	0.579	0.608	0.622
Note:			*p<0.1; **p<0.05; ***p<0.01

workers will produce enough food to meet demand, and perhaps some cocoa for export: $y_R = d_f + x_c$. If we allow food trade, however, agricultural production will equal domestic food demand plus food exports minus food imports, plus cocoa production (assumed again to be mainly for export): $y_R = d_f + x_f - m_f \rightarrow y_r + x_c = w^{\beta} - x_f - m_f \ + x_c$.

Food trade

I take net grain trade data from the International Institute of Agricultural's statistical yearbook. These are converted into calorific values, and then compared into 'net real wage equivalent' units (NRWE), by dividing by the subsistence basket calorific requirements. Thus agricultural productivity in food would be estimated as follows:

$$y_a \ = w^{\alpha} + NRWE$$

The first term on the right-hand side is the quantity of food demanded expressed in terms of a real wage subsistence basket (note that $w = w^{\alpha} = 1$ when the cost of the subsistence basket is exactly equivalent to the nominal wage). *NRWE* could be positive (the country is a net food exporter) or negative (net food importer). Adding this term ensures, for example, that high-wage net food importers like Great Britain are not credited with the production of grain that consumers demand but which must be imported. The reverse is true for net food exporters.

Non-food agricultural trade

Agricultural workers could either choose to produce food for the domestic market, to export food, or to export non-food agricultural goods like cocoa or rubber. We must therefore make allowance for this in our estimates of agricultural productivity; without it, we would assume, for example, that a country that chose to forgo food crops like millet in favour of more lucrative cash crop exports like groundnuts, and which settled its food balance with imports, was less productive than a more subsistence-oriented agricultural economy.

Operationalising this, however, is less obvious than with than with food; a conversion into caloric content makes much less sense than for food trade data. Consquently, I convert per capita non-food agricultural trade into its equivalent in silver, which I can then divide by the silver cost of the consumption basket, resulting in a quantity that is expressed in the same units as the real wage. Expressing this as 'non-staple exports' (NSE) agricultural productivity can therefore be expressed as:

$$y_a \ = w^{\alpha} + NRWE \ + \ NSE$$

This is substituted into Equation 3 instead of the simple agricultural production equation based on food demand.

Data

The resulting model requires the following pieces of data for each country:

- urbanisation rates, used to calculate the employment share of agriculture
- real wages
- staple food imports and exports, converted to calorific equivalent
- non-staple food exports
- fuel and mineral share of exports (for the calculation of the agricultural productivity gap)

The last requirement is useful but not necessary: the regression model that includes it has slightly more explanatory power over contemporary agricultural productivity gaps than those that do not, but the improvement is only modest. However, since it is likely that data sources used to construct the agricultural trade figures necessary for the rest of the model also list fuel and mineral exports, it has been included here.

For urbanisation rates, I have drawn on ClioInfra's urbanisation rate database. I have supplemented it with secondary literature for a small number of countries, like Russia and the Dutch East Indies. Real wages were constructed for another paper (Weber, Semeniuk and Westland, 2021), but they are broadly-speaking grain wages, adjusted for calorific content in the vein of (Allen, 2001). They do not always refer to 1900, but given the relative stability of wages from year to year, I consider their use broadly acceptable. Though the literature generally looks for unskilled urban wages, I sometimes found only agricultural wages, which have been inflated by 30% to account for a fairly general historical gap between rural and urban incomes. Generally speaking, a 'subsistence' amount for an adult male of each grain, tuber or other starch was calculated on the basis of calorific content: for example, 184 kilograms of rice per adult male per year. This price of this amount of food was then mutliplied by 5 to reflect the consumption—food and non-food—of a family of two adults and two children. The resulting welfare ratios are designed to be broadly comparable with Allen-style welfare ratios. Using only staple prices does introduce some degree of imprecision in the estimates, and a future extension of the present work might consider using more detailed real wage figures to improve the reliability of the resulting GDP per capita figures. As mentioned above, food balance figures were taken from the International Institute of Agriculture's annual Statistical Yearbook for 1910, the first year available; however, national trade statistics could also have been used. Non-food rural exports were taken from Weber, Semeniuk, Westland and Liang (2021).

Assessing accuracy

The Maddison database gives us 40 estimates of per capita GDP in 1901, mainly though not exclusively from Western countries, and this provides us with a set of data against which we can measure our own estimates. Taking the data described in the previous section for the subset of countries for which there is a Maddison estimate in the year 1900, I calculate GDP estimates for all countries for which we have wage and trade data, using Maddison estimates as anchors.

I have divided the countries into 5 regions: Europe and Australasia, Asia, Africa, North America, and South and Central America. I then randomly draw a GDP figure from each group to serve as the anchor country for each group. I exclude Switzerland from the European group entirely because the Maddison figures are a dramatic overestimate of actual Swiss GDP per capita in 1900 (Stohr, 2016). Because South Africa is the only country in the African region for which Maddison supplies an estimate, it is excluded from the following exercise altogether.

Relative values are computed for all of the necessary variables, using the appropriate regional anchor, and Equation 4 is estimated, giving a GDP relativity that is then applied to Maddison's estimate for the anchor country. This produces an estimate of GDP per capita for the non-anchor country. I perform this operation 500 times, with a new anchor country randomly selected each time, and calculate the median GDP per capita across all simulations, taking care to remove cases where an anchor country has been compared with itself (which results in an exact match between estimated and Maddison GDP figures). This procedure ensures that each estimate is not overly dependent on Maddison's estimates of any single anchor country.

In Table 2, I calculate perform this simulations for various versions of the method outlined above in Section , and report the root mean squared error for all GDP per capita estimates when compared to the Maddison figures. In the first set of simulations, I assume that the agricultural productivity gap is a fixed constant, equal to 3, and make no allowance for unbalanced agricultural trade. Agricultural production is simply equal to food demand, which is proportional to the square root of the real wage. In the second set, I allow p to vary, estimating it separately for each country based on the coefficients for regression (4) in Table 1—however, agricultural trade is still assumed to balance. In the third set, I keep the country-varying p parameter, but add an adjustment for staple food trade, as outlined in subsection . Finally, in the fourth set of simulations, I add an adjustment for non-staple food exports (for example, cocoa or tea exports). Importantly, the root mean-squared error declines with each additional adjustment. Though an average error of $553 seems reasonably large, this is driven by several major outliers; using the median rather than the mean squared error results in an average error of $366.

Table 2: Predictors of the agricultural productivity gap

Model	RMSE
p=3, no rural trade	$1583
varying p, no rural trade	$975
varying p, food trade only	$844
varying p, all rural trade	$533

Comparing the absolute deviation between the Maddison figures is useful, but the proportional difference is also of interest. The mean absolute proportional error is 16.8% for the fixed-*p* version of the model, but below 1% for the three other models that allow for a varying *p*. In Figure 3, I plot my estimates of per capita GDP, using the fourth 'model' with varying *p*'s and rural trade accounted for, (on the y-axis) against the Maddison Project's (on the x-axis). Both are expressed in terms of 2010 Gheary-Khamis dollars. The correlation is very high, though with some outliers (correlation coefficient of 0.96).

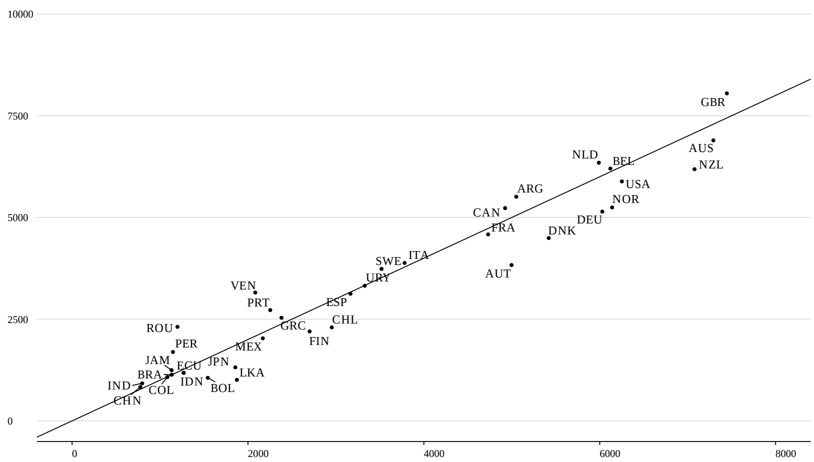

Figure 3: Maddison Project (x-axis) and this paper's (y-axis) estimates of per capita GDP c.1900.

Estimating non-Maddison countries

The real interest of devising this shortcut method lies in expanding the Maddison database. Because our trade and wage data covers more countries than the Maddison database provides GDP estimates for, we can expand considerably the number of per capita GDP estimates for the early twentieth century. I follow the same procedure as above, (with the obvious exception that anchor countries are restricted to only those countries with Maddison GDP estimates in 1900; because South Africa is the only African country in the Maddison database for this period, it is the anchor country for all African countries in every simulation). In the table below, I give per capita GDP estimates for all countries in our dataset.

Table 3: GDP per capita estimates for countries not in Maddison database in 1900 in Gheary-Khamis dollars

Country	GDP estimate	Country	GDP estimate
Angola	853	Iceland	3214
Antigua	742	Israel	2413
Benin	1355	Kenya	734
Bulgaria	2277	Lebanon	1777
Bahamas	1583	Lesotho	539
Belize	1571	Morocco	1701
Barbados	928	Madagascar	692
Côte d'Ivoire	633	Myanmar	1686
Cameroon	562	Mozambique	496
Belgian Congo	509	Mauritius	2309
French Congo	605	Malawi	425
Costa Rica	2379	Nigeria	1396
Cyprus	1722	Nicaragua	987
Djibouti	1052	Nepal	595
Algeria	1928	Philippines	1341
Ethiopia	824	Senegal	1709
Ghana	1494	Sierra Leone	979
French Guinea	897	Serbia	1890
Gambia	1610	Trinidad and Tobago	1876
Grenada	971	Tunisia	1665
Guatemala	1463	Turkey	1355
Guyana	1869	Tanzania	641
Iran	2194	Uganda	702
Iraq	2280	Vietnam	1010

Agriculture and the world economy, c.1900-2000

One of the advantages of these new estimates is that there is an explicit sectoral breakdown available, though in some cases the estimates seem questionable. These can be compared with more modern figures to chart the decline of agriculture's share of economic output across the twentieth century for most countries. In Figure 4 I plot this trajectory for ten randomly selected countries. Nine out of ten countries experienced a decline in the agricultural share of GDP; Burma was the only country to buck this trend. The most dramatic declines tend to be for African countries, some of which were almost entirely agricultural in 1900 and are now much more urbanised and their economies service-oriented.

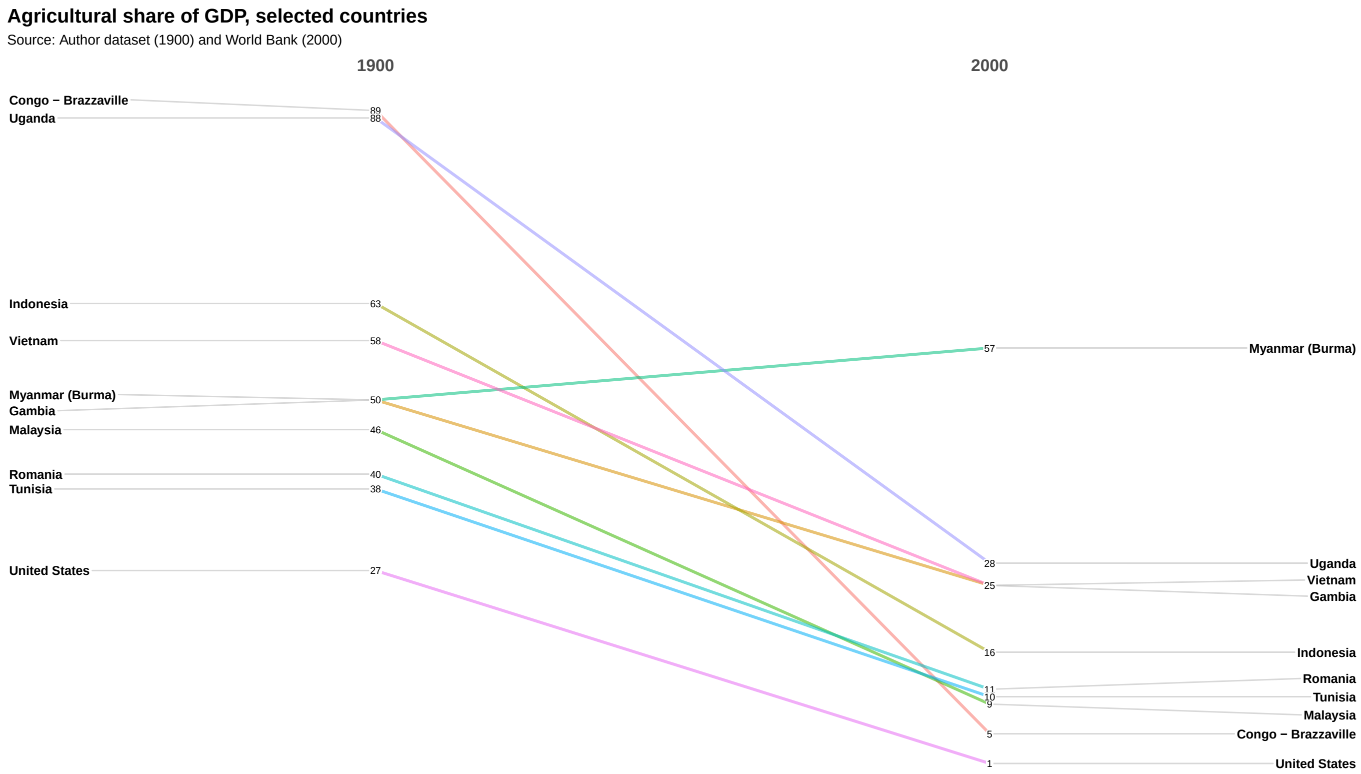

Figure 4: Agriculture's share of output, 1900-2000, selected countries.

References

Allen, R.C., 2001. The great divergence in European wages and prices from the Middle Ages to the First World War. Explorations in economic history 38, 411–447.

Bolt, J., Inklaar, R., de Jong, H., van Zanden, J., 2018. Rebasing Maddison: new income comparisons and the shape of long-run economic development. Groningen Growth and Development Centre, University of Groningen.

Gollin, D., Lagakos, D., Waugh, M.E., 2014. The agricultural productivity gap. The Quarterly Journal of Economics 129, 939–993.

Malanima, P., 2011. The long decline of a leading economy: GDP in central and northern Italy, 1300–1913. European Review of Economic History 15, 169–219.

Nunn, N., Puga, D., 2012. Ruggedness: The blessing of bad geography in Africa. Review of Economics and Statistics 94, 20–36.

Stohr, C., 2016. Trading gains: new estimates of Swiss GDP, 1851 to 2008. Working Papers of the Paul Bairoch Institute of Economic History.