Understanding per-capita income growth in preindustrial Europe

Nils-Petter Lagerlöf*
www.nippelagerlof.com
Email: lagerlof@yorku.ca

August 2016

Abstract: Fouquet and Broadberry (Journal of Economic Perspectives 2015) have recently compiled detailed time-series data over per-capita incomes for several European countries from as early as 1300, up to 1800. The time series are all volatile and highly persistent; per-capita incomes move in decades-long cycles of expansions and contractions. The current paper examines a Malthusian model with realistic life-cycle structure and stochastic growth in agricultural productivity. This model can generate per-capita income dynamics similar to the Fouquet-Broadberry data.

*I am grateful for comments from Carl-Johan Dalgaard, David de la Croix, James Fenske, Oded Galor, Marc Klemp, Anastasia Litina, Ömer Özak, Nuno Palma, Holger Strulik, and attendants at my presentations at a 2016 Macro Lunch at Brown University, the 1st CREA Workshop on Aging, Culture, and Comparative Development at the University of Luxembourg in 2016, and the 2016 SED Meetings in Toulouse. I am also grateful for financial support from the Social Sciences and Humanities Research Council of Canada. All errors are mine.

1 Introduction

A common understanding of Malthusian models is that they predict per-capita incomes to be independent of land productivity. In a Malthusian world, if one economy has higher agricultural productivity than another, they should still have the same standards of living, as long as they are otherwise identical. Changes over time in land productivity should not translate into changes in living standards.

This contrasts with what we see in the data for several supposedly Malthusian societies. Fouquet and Broadberry (2015) have used various sources to compile annual per-capita GDP data for a handful of European countries from as early as 1300, up to 1800. These all display big fluctuations in per-capita GDP over time, and on average the levels seem to trend upwards (see Figure 1).

However, those predictions refer to steady state outcomes in a non-stochastic environment, where levels or growth rates of land productivity are constant. If Malthusian economies are subject to productivity shocks, then per-capita incomes always differ from the steady-state level associated with a non-stochastic version of the model. In fact, as this paper tries to demonstrate, in a stochastic Malthusian environment, per-capita incomes can differ a great deal between any two economies, and they can be highly persistent over time within economies, on an order of magnitude comparable to Fouquet and Broadberry's data. Moreover, if *rates* of productivity growth trend upward over time, so will *levels* of per-capita incomes.

To make this point, we examine a Malthusian model with realistic life-cycle structure: each period represents one year, permitting comparison to annual data. There is growth in land productivity, and that growth rate is stochastic and its mean is increasing over time.

The model otherwise relies on the standard Malthusian building blocks. There is only one sector, producing a good that we can interpret as food. When a land productivity shock raises per-capita incomes, fertility increases and mortality falls, leading to a population expansion. Because land is in fixed supply per-capita incomes must subsequently decline. The model is thus by construction unable to generate *sustained* growth in per-capita incomes.

We then simulate 1000 model economies for 501 years, representing 1300-1800, for plausible parameter values. For each simulated economy we measure the mean, standard deviation, and serial correlation in per-capita incomes over the 501 years, and compare these to the Fouquet-Broadberry data.

Qualitatively, the model-generated paths are strikingly similar to those in the data, with several expansions and contractions lasting more than a century. Quantitatively, gaps in per-capita incomes, and levels of persistence and standard deviation over half a millennium, are all similar to what we see in the data.

The results hinge on a few assumptions that all seem realistic. First, part (but not all) of the productivity shocks must be persistent. That is, the *growth rate* of productivity, not only its level, must be subject to serially independent shocks. This makes sense in this context. Consider, e.g., increases in land productivity following the introduction of New World crops, like the potato: once introduced it did not go away. This translates to a permanent rise in productivity levels, or a one-time shock to the growth rate.

Another important assumption, if the model is to produce the (seemingly non-Malthusian) upward trend in per-capita incomes, is that the expected productivity growth rate increases over time. This assumption is also realistic, because population growth rates increased over this period. Indeed, we set productivity growth rates to match the simulated population levels to data.[1]

Finally, the elasticities of fertility and mortality with respect to wages must not be too large. Intuitively, low elasticities amplify and prolong the effects of productivity shocks, by reducing the speed at which population levels adjust. But again, empirical studies do indicate elasticities low enough for the model to match the data.

This paper relates to a literature arguing that the origin of today's world income distribution can be traced in one way or another to the timing of the transition of a few economies in Western Europe out of what is typically labeled Malthusian stagnation (e.g., Galor and Weil 2000, Hansen and Prescott 2002, Lucas 2002, Galor 2008). Critics have pointed to the non-stagnant preindustrial environment in the Fouquet-Broadberry data as evidence against such theories; other than Fouquet and Broadberry (2015) themselves, see, e.g., Dutta et al. (2016). It is in the context of that debate that this exercise becomes so important.

It is also true that variations on the standard Malthusian model can indeed alter its prediction about stagnant living standards. For example, in settings with multiple goods, and not only food, certain shocks can raise living standards permanently (Sharp et al. 2012, Voigtlander and Voth 2013, Dutta et al. 2016). Other variations on the Malthusian framework can generate endogenous fluctuations over time in living standards, such as Dalgaard and Strulik (2015), who allow for endogenous investments in body mass and thus subsistence requirements.[2]

To be clear, these extensions of the Malthusian model are important, and can be motivated independently of the lack of stagnation in the Fouquet-Broadberry data. However, the current exercise can shed light on whether such extensions are needed to account for the patterns observed in the data.

Our exercise also relates to papers applying realistic life-cycle dynamics to overlapping-

[1] Rising growth rates in productivity during the Malthusian era is also a feature of many unified growth models. For example, in Galor and Weil (2000) the rate of change in technology is an increasing function of population levels at the Malthusian stage of development.

[2] See the simulation of the Galor-Weil model by Lagerlöf (2006) for another example of endogenous cycles in a Malthusian context.

generations frameworks, often in continuous time (e.g., Boucekkine et al. 2002; de la Croix and Licandro 2013). These models are typically non-stochastic, but share some of the mechanisms through which shocks can propagate themselves in our simulations.

Several papers have calibrated growth models to similar data, often with an aim to produce a transition from stagnation to growth in per-capita incomes (e.g., Fernández-Villaverde 2001; Lagerlöf 2003, 2006; Bar and Leukhina 2010). This paper is probably the first to examine whether a purely Malthusian model, unable to generate sustained growth by construction, can generate transition-like time paths when adding shocks to land productivity.

Finally, this paper is motivated by other evidence that the preindustrial world was Malthusian. Across regions defined by modern country borders, preindustrial levels of technology correlate more strongly and robustly with population densities than with living standards (Ashraf and Galor 2011). Moreover, the very same preindustrial economies displaying non-stagnant incomes in the Fouquet-Broadberry data seem to react in a Malthusian fashion to exogenous changes in per-capita incomes. Across Swedish counties in the 19th century, good harvests were associated with higher birth rates and lower death rates (Lagerlöf 2015). Similar patterns have been found for other Scandinavian countries in time-series data (Klemp and Møller 2016), and for England (Nicolini 2007, Crafts and Mills 2009, Kelly and Ó Gráda 2014, Klemp and Møller 2016).

The rest of this paper is organized as follows. Section 2 summarizes some of the facts about per-capita incomes in preindustrial Europe that we learn from Fouquet and Broadberry (2015) and their sources. Section 3 sets up the model, first illustrating the main mechanisms in a stylized version in Section 3.1, and then presenting a setting realistic enough to be simulated in Section 3.2. Section 4 presents results from the simulation and compares them to the data. Section 5 concludes.

2 The facts

Figure 1 shows log per-capita GDP for five economies (with periods indicated in parenthesis): England/UK (1300-1800), Italy (1310-1800), the Netherlands (1348-1800), Portugal (1530-1800), and Sweden (1560-1800). The data are here reported in logs and normalized so that the logged series equal zero when averaged over time and across countries.

The source for this data is Fouquet and Broadberry (2015), who in turn rely on various in-depth studies for the respective countries.[3] They also report data from Spain but not on

[3] The original sources are as follows: for England/UK: Broadberry et al. (2011); for Italy (more precisely its central and northern parts): Malanima (2011); for the Netherlands (more precisely the province of Holland): van Zanden and van Leuwen (2012); for Sweden: Schön and Krantz (2012); for Portugal: Palma and Reis (2016). Fouquet and Broadberry (2015) also cite a paper by Reis, Martins, and Costa as source

annual frequency so we do not use those numbers here.

Four things can be noted from Figure 1. First, per-capita GDP was not constant over this period, but fluctuated a lot for all countries. Any model that seeks to replicate such time series arguably needs some stochastic component.

Second, because of these fluctuations, there are noticeable differences in GDP per capita across these countries in any given year, and some of these gaps stay when averaged over time. Portugal is poorest on average, and the Netherlands richest; cf. Table 5.

Third, average log GDP per capita across the five countries shows a mild upward trend over time. This is easiest to see for the period after 1560 when data is available for all five countries, but the same holds for earlier periods when considering only countries with available data. In short, levels for most countries grow a little over time.

Fourth, the time series show a great deal of persistence. A country's level of per-capita GDP in any given year is highly correlated with where it was in the previous year, and even two or three years back. Put another way, GDP per capita expands and contracts in long cycles, rather than jumping all over the place from one year to another.

The next section discusses how these patterns could be reconciled with a Malthusian model.

3 Theoretical framework

3.1 A simplified setting

To compare any annual data to model-generated time paths it helps to let each model period correspond to one year. Section 3.2 below considers such a model. However, it is useful to first illustrate some of the mechanisms driving the results using an overlapping-generations framework where agents live for only two periods. In the first phase of life, they are inactive children; in the second, adult, phase, they earn income, consume, and rear children. To fix notation, let agents who are adult in period t earn wage w_t, consume c_t, and rear n_t children.

There are two twists to the framework presented here, compared to most textbook Malthusian models: (1) an income-fertility elasticity less than one, and (2) sustained (but for the moment non-stochastic) growth in land productivity.

To capture the first of these model innovations, let the cost of rearing n_t children be $qn_t^{1/\delta}$, where $q > 0$ and $\delta \in (0, 1]$. That is, δ is the inverse of the elasticity of the child-cost function, which will soon be seen to determine the elasticity of optimal fertility with respect to wages. Most standard Malthusian models assume $\delta = 1$.

for Portugal, but the actual data is from Palma and Reis (2016); I thank Nuno Palma for pointing this out to me.

The budget constraint can now be written

$$c_t = w_t - q n_t^{\frac{1}{\delta}}. \tag{1}$$

Utility is logarithmic and defined over the number of children, and (adult) consumption, with weight $\widetilde{\gamma} \in (0,1)$ on the former:

$$U_t = (1 - \widetilde{\gamma}) \ln(c_t) + \widetilde{\gamma} \ln(n_t). \tag{2}$$

Maximizing (2) subject to (1), some algebra gives the agent's optimal fertility as follows:[4]

$$n_t = \gamma w_t^{\delta}, \tag{3}$$

where

$$\gamma = \left(\frac{\delta \widetilde{\gamma}}{q \left[1 - \widetilde{\gamma}(1 - \delta)\right]} \right)^{\delta}. \tag{4}$$

That is, δ measures the elasticity of fertility with respect to wages.

Total output in period t equals

$$Y_t = (M A_t)^{\alpha} L_t^{1-\alpha}, \tag{5}$$

where α is the land share of output, M is total land size, A_t is a land-augmenting productivity factor, and L_t is the size of the labor force, which is the same as the adult population.

Labor is paid its marginal product.[5] Using (5), and normalizing land size to unity ($M = 1$), the wage rate can thus be written

$$w_t = (1 - \alpha) \left(\frac{A_t}{L_t} \right)^{\alpha}. \tag{6}$$

The second novelty compared to most other Malthusian models is the assumption of sustained growth in land productivity, here set to some exogenous and constant rate g:

$$A_{t+1} = (1 + g) A_t. \tag{7}$$

Each adult agent has n_t children—all of whom are assumed (for now) to survive until adulthood—and since all agents die after the adult phase of life, the labor force evolves according to $L_{t+1} = n_t L_t$. Forwarding (6) to period $t+1$, and applying (3), (7), and

[4] Section A of the Appendix shows how to derive (10), which is the corresponding relationship in the extended model set up in Section 3.2.

[5] Output not paid to labor is here implicitly assumed to be allocated to a landowning elite of small fixed population size, playing no role in the rest of the analysis.

$L_{t+1} = n_t L_t$, gives a first-order difference equation for the wage rate:[6]

$$w_{t+1} = \left(\frac{1+g}{\gamma}\right)^{\alpha} w_t^{1-\alpha\delta}, \tag{8}$$

which has a unique and stable steady-state equilibrium, defined by

$$\overline{w} = \left(\frac{1+g}{\gamma}\right)^{\frac{1}{\delta}}. \tag{9}$$

Two qualitatively important insights can be gained from (9). First, the steady-state wage rate, $\overline{w}$, is higher in economies with faster productivity growth, higher g. This is quite intuitive. When the wage rate is constant at $\overline{w}$, the ratio A_t/L_t is also constant; recall (6). Thus, population grows at the same (gross) rate as land productivity, $n_t = 1+g$, which from (3) requires higher w_t in steady state. In other words, for population to keep up with faster productivity growth, living standards must be higher.

The second insight from (9) is a corollary of the first. The positive effect on $\overline{w}$ from increasing g is stronger when δ is small. Intuitively, the less elastic is fertility to changes in wages, the more wages must increase in response to an increase in productivity growth to keep productivity-population ratio constant. If δ is small—and it seems to be at least less than one in the data—then small differences in g can generate large gaps $\overline{w}$.

This may explain how two otherwise identical Malthusian economies can have different living standards, with relatively modest differences in productivity growth rates. Moreover, patterns like those in Figure 1 could possibly occur if growth rates fluctuate over time. To explore this possibility requires that these mechanisms are nested in a less stylized framework, a task undertaken in the next section.

3.2 A more realistic setting

3.2.1 Fertility

Consider now an overlapping-generations model where agents live for at most T periods. They are reproductive from period $\underline{B}$ to $\overline{B}$, and earn wages from period $\underline{B}$ to R, where $1 < \underline{B} < \overline{B} < R \leq T$.

In each model period t, an agent in period $j \in \{\underline{B}, ..., R\}$ of life earns a wage $w_{j,t}$, which determines the number of children conceived in that period (born in the next), analogously

[6] To see this, note that

$$w_{t+1} = (1-\alpha)\left(\frac{A_{t+1}}{L_{t+1}}\right)^{\alpha} = (1-\alpha)\left(\frac{A_t}{L_t}\right)^{\alpha}\left(\frac{1+g}{\gamma w_t^{\delta}}\right)^{\alpha},$$

and recall (6).

to (3):

$$n_{j,t} = \gamma_j w_{j,t}^{\delta}, \tag{10}$$

where $\delta > 0$ is the elasticity of fertility with respect to wages, and $\gamma_j > 0$ is here an age-specific parameter, such that $\gamma_j = 0$ for $j \notin \{\underline{B}, ..., \overline{B}\}$.

Section A of the Appendix presents a simple model, which generates the behavior postulated in (10), as well as consumption of all agents (including those who do not earn incomes). For the purpose of the current modeling exercise, we do not need to know where the behavior described in (10) comes from.

3.2.2 Production and land productivity

In any period t, total output of a single good, Y_t, is produced using a Cobb-Douglas production function, with land and effective labor as inputs, similar to the model in Section 3.1:

$$Y_t = (X_t A_t)^{\alpha} L_t^{1-\alpha}, \tag{11}$$

where L_t is effective labor (explained further below), α is the land share of output, and the amount of land is (again) normalized to unity. Land productivity in period t equals $X_t A_t$, where X_t and A_t are subject to temporary and permanent shocks, respectively. These shocks are distributed as follows:

$$\ln(X_t) \sim N(0, \sigma_X), \tag{12}$$

and

$$\ln(A_{t+1}) - \ln(A_t) \sim N(\mu_t, \sigma_A), \tag{13}$$

where μ_t is the expected productivity growth rate, which is assumed to be time-dependent (but non-stochastic), allowing mean growth in land productivity to change over time. The parameters σ_X and σ_A denote the standard deviations in the temporary and permanent shocks, respectively. Temporary shocks could represent fluctuations in weather. Permanent shocks could capture innovations to agricultural technology, or the effects of newly introduced crops; see Nunn and Qian (2011) for evidence of a positive effect on population levels from the introduction of potato.

In any model period t, let $P_{j,t}$ be the population in the jth period of life, and let each age group $j \in \{\underline{B}, ..., R\}$ supply one unit of labor. Effective labor is determined by a CES aggregation function, allowing flexible substitutability between labor inputs of the different working age groups:

$$L_t = \left[\sum_{j=\underline{B}}^{R} \beta_j P_{j,t}^{\rho} \right]^{\frac{1}{\rho}}, \tag{14}$$

where $\rho < 1$, and $\sum_{j=\underline{B}}^{R} \beta_j = 1$. The most standard assumption might be that $\rho = 1$, and that β_j is constant across age groups, implying perfect substitutability between labor supply of different cohorts. The alternative assumption ($\rho < 1$, and β_j being different across cohorts) can be interpreted as different age cohorts performing different work tasks, perhaps because older workers have more experience, and younger workers more physical strength. This allows the age distribution in any given period to have an effect on output, and thus incomes, and reproduction.

3.2.3 Wages

Recall that an agent in period j of life earns wage $w_{j,t}$, which equals the marginal product of that age group's labor input. Using (11) and (14), some algebra shows that

$$w_{j,t} = \frac{\partial Y_t}{\partial L_t}\frac{\partial L_t}{\partial P_{j,t}} = (1-\alpha)\frac{Y_t}{L_t}\beta_j\left(\frac{L_t}{P_{j,t}}\right)^{1-\rho}. \tag{15}$$

Using (14) and (15), it can be seen that total payments to workers can be written

$$\sum_{j=\underline{B}}^{R} w_{j,t}P_{j,t} = (1-\alpha)Y_t. \tag{16}$$

Although it does not matter for the analysis, we here implicitly assume that the remaining output, αY_t, is allocated to landowners, who we can think of as old and non-active agents, using their income for consumption (see Section A of the Appendix).[7]

3.2.4 Population dynamics

The new-born population in period $t+1$, $P_{1,t+1}$, is made up of children conceived in period t. A reproductive agent in the jth period of life conceives $n_{j,t}$ children, and since there are $P_{j,t}$ in each such cohort, it follows that

$$P_{1,t+1} = \sum_{j=\underline{B}}^{\overline{B}} n_{j,t}P_{j,t}. \tag{17}$$

For all other cohorts, population levels evolve according to

$$P_{j+1,t+1} = s_{j,t}P_{j,t}, \tag{18}$$

where $s_{j,t} \in [0,1]$ denotes the rate at which agents survive from the jth period of life to the next (and from model period t to the next). We describe $s_{j,t}$ further below.

[7] Alternatively, land income could be allocated to a class of elite agents of small and fixed size, rearing only one offspring per agent, and using the reminder of their income for consumption.

3.2.5 Total population and per-capita incomes

Let total population be denoted

$$P_t = \sum_{j=1}^{T} P_{j,t}. \tag{19}$$

For any given levels of output, Y_t, and total population, P_t, the economy-wide per-capita income level becomes

$$y_t = \frac{Y_t}{P_t}. \tag{20}$$

In what follows, we shall let y_t in the model correspond to GDP per capita in the data.

3.2.6 The survival rate

The model allows for three mortality factors: starvation, disease, and age. We want to incorporate each of these, since they may all to some extent have impacted per-capita income dynamics in the European data. For example, disease shocks could be one factor contributing the volatility in per-capita incomes. To that end, the survival rate is defined as

$$s_{j,t} = s_t^y s_t^d s_j^{\text{age}}, \tag{21}$$

where s_t^y, s_t^d, and s_j^{age} all lie on $[0,1]$, and represent survival from starvation, disease, and age, respectively. Survival from starvation (s_t^y) is specified as

$$s_t^y = \min\left\{1, \left(\frac{y_t}{\overline{y}}\right)^{\kappa}\right\} \in (0,1], \tag{22}$$

where y_t is per-capita output in (20), $\overline{y}$ is the corresponding (non-stochastic, non-growing) steady-state level of y_t, and κ is a parameter measuring the elasticity of the survival rate with respect to falls in per-capita incomes. In this formulation, the mortality effects are present only when living standards are sufficiently low. This is broadly consistent with data from 19th-century Sweden, where good harvests do not lower mortality much, but bad harvests raise mortality (Lagerlöf 2015). It also seems intuitive that mortality from malnutrition is constrained to zero when food intake is above some threshold level. Here that threshold is set at the level associated with a Malthusian steady state absent shocks and starvation.[8]

The disease component in (21) is defined as

$$s_t^d = \exp\left(-\phi m_t^2\right) \in (0,1], \tag{23}$$

where $m_t \sim N(0,1)$ is a mortality shock, such that larger deviations of m_t from its zero mean imply lower survival rates, and ϕ is a parameter capturing the size of the effect of

[8] To be precise, we let $\overline{y}$ be given by the level derived in Section C of the Appendix, which applies to the case with unit elasticity of fertility ($\delta = 1$), enabling us to easily find an analytical expression for $\overline{y}$.

these shocks. It can be shown that $E(s_t^d) = (1+2\phi)^{-.5} \approx 1-\phi$, so ϕ is approximately the expected annual death rate from disease.

Finally, the age component in (21), s_j^{age}, varies with age, but is constant over time. It is set to match data from Sweden, as explained below.

4 Quantitative analysis

4.1 Benchmark parameter values

To generate simulated time paths to compare to data we first need to make assumptions about parameter values. Most of these are summed up and explained in Table 1.

First, the life-cycle parameters $\underline{B}$, $\overline{B}$, R and T can be set very intuitively. Agents live for at most $T = 90$ years, reproduce between life periods $\underline{B} = 15$ and $\overline{B} = 49$, and work from $\underline{B} = 15$ to $R = 70$, all conditional on survival.[9]

The land share of output (α) is set to 0.4, as in Hansen and Prescott (2002).

The reproduction parameters (γ_j), the age-dependent weights in the production function (β_j), and the elasticity of substitution between labor inputs of different age groups (ρ), are set to jointly match age-specific fertility and wage data from Sweden. Fertility data are averages from 1750-1800, and the wage data refer to agricultural workers in 1940 (the earliest year for which such age-wage data is available, but Sweden was at the time particularly dependent on its agricultural sector due to the second world war); see Section B of the Appendix for further details. Values of γ_j and β_j are shown in Table 3 for three age groups. The model-data match is illustrated in Figure 2.

Consider next the elasticity parameters, δ and κ. Klemp and Møller (2016, Appendix B) provide a summary of estimated elasticities of fertility and mortality with respect to wages. These are so-called long-run elasticities (sums of elasticities with respect to wages at various time lags), and based on aggregated time-series data over Crude Birth and Death Rates (i.e., total births and deaths over total population), mostly from England and starting as early as 1540. Estimates of fertility elasticities range from .12 to .32, and mortality elasticities from $-.47$ to $-.08$; most studies find numbers at the lower end of those intervals.

Since many of these are based on aggregate data one may worry about causality; part of the variation could be driven by economic activity declining in periods when agents reduce child rearing for other reasons than direct Malthusian checks, e.g., disease or war. Partly addressing such concerns, Lagerlöf (2015) exploit cross-county harvest fluctuations in Sweden

[9] To be precise, setting $\overline{B} = 49$ means that, in the last period of life in which an agent can *conceive* a child, the agent's *age* is 48; the agent is thus in the 49th year of life. The child is born in the 50th period of the parent's life, when the parent is of age 49. Similarly, agents can start to conceive (a small but positive number of) children when they are of age 14 (in the 15th year of life).

1816-1856. Those elasticity estimates are somewhat smaller: .1 and $-.09$ for fertility and mortality, respectively.

Section B.4 of the Appendix describes in more detail how to compare the model-generated survival and fertility rates to these estimates. In short, we compute the CBR and CDR from our simulated data. We then regress the logarithm of these variables on logged wages at various lags; see Table A.1 in the Appendix. We set $\delta = .15$ and $\kappa = .01$, putting the elasticities estimated from the simulated data at around .13 and $-.2$, for birth and death rates, respectively. These fall within the ranges of existing empirical estimates.

We set $\phi = .01$, implying that the expected annual death rate from disease is about 1%.

Finally, s_j^{age} is set to make $E(s_t^d)s_j^{\text{age}} \approx (1-\phi)\, s_j^{\text{age}}$ roughly fit mortality data by age from Sweden 1750-1800, with s_{90}^{age} set to zero, so that all agents who survived for $T = 90$ periods die with certainty after that; see Table 2.

The time-dependent expected growth rate in land productivity, μ_t, is set to increase from 0% per year in 1300 to 1.25% in 1800, with more of the increase coming toward the end of that period. This makes the model match the double exponential trends in population levels for the five countries, as shown in Figure 3, and also implies a rising trend in levels of GDP per capita in the simulation.

The parameters measuring the dispersion in the shocks in (12) and (13), σ_X and σ_A, are set so that the simulations generate moments of $\ln y_t$ measured over 501 years as close as possible to (or with at least some overlap with) the corresponding numbers in the data. This is explained in connection to the results in Figure 6 below.

4.2 Initial conditions

We need to set start values so that the initial distributions are close to the steady-state distributions to which they would eventually converge absent the trend in expected productivity growth. Otherwise, the model might artificially generate high persistence, due to a standard convergence process over the first several periods. To that end, we set initial values at the steady-state levels associated with a deterministic version of the model, with unit fertility elasticity; analytical expressions for these initial values are derived in Section C of the Appendix.

We then simulate the model (without productivity growth, $\mu_t = 0$, but with all other parameters set as in the benchmark case above) over 500 periods before starting to measure outcomes. The 501st simulated period represents the year 1300.

4.3 Simulation results

Among 1000 simulated economies under our benchmark parametrization, Figure 4 shows the time paths of $\ln y_t$ for the first four; recall that y_t in (20) corresponds to GDP per capita in the data. The paths are normalized to equal zero when averaged over time and across simulated economies, similar to the data in Figure 1. As seen, in any given year we observe large gaps in $\ln y_t$ across these four economies, and each of them displays long cycles of expansions and contractions over time. The patterns are strikingly similar to those based on actual data in Figure 1. All the economies are parametrically identical, with all shocks drawn from the same distributions, so the differences between them are driven only by different realizations of the shocks.

Recall also that the model which generates the paths is stagnant by construction, in the sense that per-capita incomes cannot exhibit sustained growth; each of the growth spurts in Figure 4 is eventually followed by a decline. Not knowing this, one could easily interpret some of the growth phases as break-outs from stagnation.

The mean of $\ln y_t$ across all 1000 runs shows a mild upward trend in Figure 4. This is generated by the upward drift in productivity growth rates, as captured by the rise in μ_t. To understand why, recall the simplified Malthusian model in Section 3.1, where a higher productivity growth rate is associated with higher per-capita incomes (and wages) on the balanced growth path; cf. (9). When this growth rate increases gradually over time the result is a gradual rise in per-capita incomes.

While Figure 4 shows just a few random paths, Figure 5 shows the 5th and 95th percentiles of $\ln y_t$ in any given year among the 1000 simulated economies. That is, in any given year, 90 percent of the simulated economies fall between these two percentiles. The corresponding paths for the five countries in Figure 1 are also shown. The paths sometimes fall outside of the interval, but mostly within. As shown in Table 4, there is variation across the countries, but roughly 5% of all the country-years in the data fall below the 5th percentile, and about 3% above the 95th.

Figure 6 gives a different picture of how well the model can match the data, by displaying histograms over some time-series moments of $\ln y_t$: mean, standard deviation, and serial correlation coefficients at one- and two-year lags, each calculated over the 501-year period representing 1300-1800. The figure shows the corresponding numbers in the data as well, also displayed in Table 5, together with the 5th and 95th percentiles of the histograms from the simulations.

Consider first the top-left panel of Figure 6, which shows the distribution of means. Reflecting the different realizations of the shocks, some of the simulated economies have higher means in $\ln y_t$ than others. In principle, we can generate any amount of dispersion with large enough standard deviations of the productivity shocks, σ_X and σ_A. But if we set these too large it becomes difficult for the model to match the corresponding distribution for standard deviations of $\ln y_t$ across the 501 years, as illustrated in the top-right panel of Figure 6.

Table 5 illustrates this with some numbers. Under the benchmark setting, the model can generate just enough variation in means to make (almost) all five economies fall within the 5th and 95th percentiles of the simulated economies (Portugal being on the border), but Italy falls slightly below the 5th percentile for standard deviations. We could make the model account for Italy's low levels of standard deviation, but then it would be harder to match the low mean for Portugal.

A similar point can be made about the one- and two-year serial correlation coefficients, as shown in the bottom two panels of Figure 6, with numbers in Table 5. In data for Sweden and Italy these two coefficients fall below the bottom 5th percentile of the simulated economies. However, the differences are not huge, when considering that correlation coefficients can range from -1 to 1; note the scales used in Figure 6. There would be less serial correlation with more dispersion in the temporary shocks (higher σ_X), and/or with less dispersion in the permanent shocks (lower σ_A), but that would worsen the fit with standard deviations and means, respectively.

Another way to compare the model to data is to examine how these moments co-vary across the 1000 simulated economies, and the five real ones. This illustrated in Figure 7. The right-hand panel shows that more serial correlation is associated with higher measures of standard deviation, both across simulations and in the data. In the model, some of the shocks to land productivity are persistent, and thus generate persistence in $\ln y_t$, making higher measured serial correlation be associated with higher measures of standard deviation. As shown in the left-hand panel, there is no similar relationship between standard deviations and means, either in the data or across simulated economies. Intuitively, shocks can be both good and bad, so more volatile paths can be associated with lower or higher incomes on average.

Of course, σ_X and σ_A were set to match these patterns, but it is far from obvious that the model should be able to do as well as it does. Recall that the differences in outcomes are the result only of variations in the realized shocks across simulations. Indeed, the fact that the match is not perfect illustrates this point. If we were to let these countries differ also in some exogenous parameter—e.g., in terms of the γ_j's, which determine long-run per-capita incomes—then it would be a trivial exercise to match model and data.

Moreover, given that the data probably also come with some measurement error, the most reasonable interpretation seems to be that they do not reject the Malthusian model.

4.4 Robustness checks

We have learned that the Malthusian model can conform relatively well with the data. What assumptions drive this result? Table 5 indicates how the 5th and 95th percentiles for the various moments change when altering some assumptions.

4.4.1 Only temporary shocks

First, we close down all permanent shocks by setting $\sigma_A = 0$, and then adjust the standard deviation in the temporary shocks by setting $\sigma_X = .35$. This serves to position the distribution of simulated standard deviations somewhere in the middle of the same measures for the five countries; see the top-right panel of Figure 8. That is, we try to make the model-data match in standard deviations as good as possible, given the constraint that all shocks are temporary.

As seen in Figure 8, there is now too little serial correlation in the model-generated distributions to match data, and the mean per-capita outcomes have far less dispersion than observed between the five countries. Permanent shocks to productivity are thus needed for the model to match the data. This is quite intuitive, since any standard Malthusian model would predict that changes to *levels* in productivity have no effect on steady-state living standards, but changes in productivity *growth rates* can, as we learned in Section 3.1.

This insight in a sense resembles that from simulating real-business cycle models, where the model's internal propagation mechanisms are often so weak that all persistence comes from the shocks fed into the model (Gogley and Nason 1995). However, Figure 8 also shows that there is some propagation force at play that generates persistence in our model, since even without any serially correlated shocks, the simulated distributions of measured serial correlation have most of the mass above zero. In other words, the Malthusian model's internal mechanisms do matter for the results.

4.4.2 Using 1560 as start year

Data coverage starts in different years for each country, the latest (Sweden) lacking data before 1560, so the moments are calculated over different periods for the five different countries. The easiest way to correct for this is to consider only the period 1560-1800, for which all five countries have data. Figure 9 illustrates the outcomes when doing this for the data and the model simulations; the associated 5th and 95th percentiles from the simulations are shown in Table 5.

The means in the simulated series now show much larger dispersion, easily overlapping with the means observed in data. This is because they are calculated on smaller samples (over shorter time periods). Compared to Figure 6, the moments for four of the countries shift slightly, since they too are calculated over a different time period. The model-data fit for standard deviations is similar to that in the benchmark case, with Italy as the marginal case. Also similar to the benchmark case, the model cannot quite account for some countries' low levels of serial correlation; that of the Netherlands is lowest of them all over this period.

However, the gaps between model and data are not huge. Note also that we here keep σ_X and σ_A at their benchmark values. The fit in Figure 9 would be better if we were to recalibrate these when using 1560 as start year.

4.4.3 Higher elasticities

Two crucial parameters are those which guide the elasticities of fertility (δ) and mortality (κ). Table 5 shows how the 5th and 95th percentiles of the simulated moments change when doubling each of these over their benchmark values: δ from .15 to .3. and κ from .01 to .02. Either of these changes shrinks the dispersion in means of per-capita incomes across simulated economies, and lowers the levels overall of measured standard deviations. Intuitively, the larger is either elasticity, the smaller are the effects of (both permanent and temporary) productivity shocks, since population levels adjust faster when wages increase, making the effects on per-capita incomes less prolonged. There is little effect on the percentiles of the distributions of the serial correlation coefficients in Table 5.

4.4.4 Higher disease mortality

The parameter ϕ guides the mortality component that we associate with the disease environment, s_t^d. A higher ϕ implies lower expected survival rates from disease, as well as higher variance in the survival rate. The last row of Table 5 shows what happens when we double ϕ from .01 to .02. Mean outcomes become more dispersed, and the distribution of standard deviations shifts up, and so do (at least marginally) the distributions of the serial correlation coefficients.

This is quite intuitive; feeding larger shocks into the model makes outcomes vary more. The increase in serial correlation can be understood from the way in which a standard Malthusian model reacts to a disease-induced drops in population levels. After an immediate increase in per-capita incomes, the economy converges gradually back to its Malthusian steady state, showing up as persistence time series. A higher ϕ implies larger such shocks, and thus more serial correlation.

5 Concluding remarks

Over the last several years researchers have compiled high-quality comprehensive per-capita income data for a couple of European countries from 1300-1800, recently summarized by Fouquet and Broadberry (2015). The data display big fluctuations in per-capita incomes over time, with a high degree of persistence, as per-capita incomes can move in long cycles of expansions and contractions. The average per-capita income across these countries even shows an upward trend. Can these observations really be consistent with the predictions of a standard Malthusian model, which says that per-capita incomes should be stagnant?

This paper proposes a simple exercise to come up with a tentative answer to this question. We set up a Malthusian model that is stagnant by construction, in the sense that it cannot exhibit sustained growth in per-capita incomes, and simulate it 1000 times to compare the time paths to Fouquet and Broadberry's data. The life-cycle structure is such that each model period corresponds to one year, enabling us to compare the results to annual data. Land productivity is subject to shocks, both temporary and persistent, and the average productivity growth rate increases slightly over time.

We try to be as informed as possible when setting parameter values: parameters guiding age-profiles for mortality, wages, and fertility are set based on data from preindustrial Sweden; elasticities of mortality and fertility with respect to wages are set to match empirical studies on English and Scandinavian preindustrial time-series and panel data; the rise over time in the mean productivity growth rate is set to make the model roughly match the accelerating pace of population growth over the period.

It turns out that the model can match the data decently, as measured by means, standard deviations and serial correlations in per-capita incomes over time, given the right amount of variation in the land productivity shocks. The simulated time paths can show centuries-long cycles of expansions and contractions in per-capita incomes. Not knowing that they are generated by a Malthusian model, one could mistakenly believe that some of the expansions constitute transitions out of Malthusian stagnation. Moreover, mean per-capita income across all 1000 simulated economies shows an upward trend, mirroring the gradual rise in the productivity growth rate, in turn set to match the accelerating population growth rate over this period.

A few assumptions are important for these results. There must be growth in land productivity, and that growth rate must be stochastic and increase over time in expectation. Fertility and mortality cannot be too elastic with respect to fluctuations in wages. Like discussed, the numbers used in the simulations are consistent with empirical estimates.

This is not proof that the Malthusian model set up here is the only one to explain preindustrial growth in Europe. Other models may match the data equally well, or better. This may include models that allow for (some modest amount of) quality-quantity trade-off in children, and/or the presence of a non-agricultural sector. The aim of this paper is to start off with a model without these features—one that is closer to how, e.g., Galor and Weil (2000) and Hansen and Prescott (2002) model the pre-transition, or Malthusian, stage of development—and assess whether it can at all be consistent with the data compiled by Fouquet and Broadberry (2015). Indeed, the model-data match is not perfect, so one could argue in favor of a *somewhat* more sophisticated model. However, it seems wrong to claim that these new and important data obviously refute the Malthusian model.

Regardless of one's preferred conclusion, this paper provides a parametric and numerical framework that can be extended and used for other applications. For example, it could serve as a starting point for constructing a unified growth framework with realistic life cycles, where existing simulations of such models tend to assume that each generation lives for only two periods of 20-30 years each. Exercises such as that are left for future research.

References

[1] Ashraf, Q., and O. Galor, 2011, Dynamics and stagnation in the Malthusian epoch, American Economic Review 101, 2003-2041

[2] Bar, M., and O. Leukhina, 2010, Demographic transition and industrial revolution: a macroeconomic investigation, Review of Economic Dynamics 13, 424-451

[3] Boucekkine, R., D. de la Croix, and O. Licandro, 2002, Vintage human capital, demographic trends, and endogenous growth, Journal of Economic Theory 104, 340–375

[4] Broadberry, S., B. Campbell, A. Klein, M. Overton, and B. van Leeuwen, 2011, British economic growth, 1270–1870: an output-based approach, mimeo. Available from: http://www.lse.ac.uk/economicHistory/pdf/Broadberry/BritishGDPLongRun16a.pdf

[5] Crafts, N., and T.C. Mills, 2009, From Malthus to Solow: how did the Malthusian economy really evolve?, Journal of Macroeconomics 31, 68-93

[6] Dalgaard, C-J., and H. Strulik, 2015, The physiological foundations of the wealth of nations, Journal of Economic Growth 20, 37-73

[7] de la Croix, D., and O. Licandro, 2013, The child is father of the man—implications for the demographic transition, The Economic Journal 123, 236-261

[8] Dutta, R., D.K. Levine, N.W. Papageorge, and L. Wu, 2016, Entertaining Malthus: bread, circuses and economic growth, mimeo, available from: http://www.dklevine.com/archive/refs4786969000000001365.pdf

[9] Fernández-Villaverde, J., 2001, Was Malthus right? Economic growth and population dynamics, mimeo, University of Pennsylvania

[10] Galor, O., 2008, The 2008 Lawrence R. Klein lecture—comparative economic development: insights from Unified Growth Theory, International Economic Review 51, 1-44

[11] Galor, O., and D. Weil, 2000, Population, technology, and growth: from Malthusian stagnation to the demographic transition and beyond, American Economic Review 90, 806-829

[12] Cogley, T. and J.M. Nason, 1995, Output dynamics in real-business-cycle models, American Economic Review 85, 492-511

[13] Hansen, G.D., and E.C. Prescott, 2002, Malthus to Solow, American Economic Review 92, 1205-1217

[14] Kelly, M., and C. Ó Gráda, 2014, Living standards and mortality since the middle ages, Economic History Review 67, 358–381

[15] Klemp, M., and N.F. Møller, 2016, Post-Malthusian dynamics in pre-industrial Scandinavia, Scandinavian Journal of Economics, forthcoming

[16] Lagerlöf, N.P., 2003, From Malthus to modern growth: can epidemics explain the three regimes?, International Economic Review 44, 755-777

[17] Lagerlöf, N.P., 2006, The Galor-Weil model revisited: a quantitative exercise, Review of Economic Dynamics 9, 116-142

[18] Lagerlöf, N.P., 2015, Malthus in Sweden, Scandinavian Journal of Economics 117, 1091-1133

[19] Lucas, R.E., 2002, Lectures on Economic Growth, Harvard University Press

[20] 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

[21] Nicolini, E.A., 2007, Was Malthus right? A VAR analysis of economic and demographic interactions in pre-industrial England, European Review of Economic History 11, 99-121

[22] Nunn, N, and N. Qian, 2011, The potato's contribution to population and urbanization: evidence from a historical experiment, Quarterly Journal of Economics 126, 593-650

[23] Palma, N. and J. Reis, 2016, From convergence to divergence: Portuguese demography and economic growth, 1500-1850, Groningen Growth and Development Centre Research Memorandum 161. www.ggdc.net/publications/memorandum/gd161.pdf

[24] Schön, L., and O. Krantz, 2012, The Swedish economy in the early modern period: constructing historical national accounts 1560–2000, European Review of Economic History 16, 529–49

[25] Sharp, P., H. Strulik, and J. Weisdorf, 2012, The determinants of income in a Malthusian equilibrium, Journal of Development Economics 97, 112-117

[26] Statistics Sweden, 1942, Lönestatistisk årsbok för Sverige 1940, Isaac Marcus boktryckeri-aktiebolag, Stockholm. Published online by Statistics Sweden: www.scb.se/Grupp/Hitta_statistik/Historisk_statistik/_Dokument/SOS/Loner/Lonestatistisk-arsbok-for-Sverige-1940.pdf

[27] Statistics Sweden, 1969, Historical Statistics of Sweden, Part 1: Population 1720-1967, AB Allmänna Förlaget, Stockholm. Available from: www.scb.se/Grupp/Hitta_statistik/Historisk_statistik/_Dokument/Historisk-statistik-for-Sverige-Del-1.pdf

[28] van Zanden, J.L., and B. van Leeuwen, 2012, Persistent but not consistent: the growth of national income in Holland 1347–1807, Explorations in Economic History 49, 119–30

[29] Voigtlander, N., and J. Voth, 2013, The three horsemen of riches: plague, war and urbanization in early modern Europe, Review of Economic Studies 80, 774-811

APPENDIX

A Micro foundations for the model in Section 3.2

This section proposes a model, similar to the simplified one in Section 3.1, that allows us to derive the fertility behavior in (10), as well as an expression describing consumption of all agents in the economy.

The model builds on a few, more or less plausible, assumptions. First, agents care about births, rather than the stock of born (and living) children. Second, there are no means of saving. These first two assumptions allow us to abstract from intertemporal allocation decisions over long time horizons.

Third, children carry only goods costs, and no time costs. This is probably not important, but simplifies the analysis, since working agents can then be assumed to supply labor inelastically.

Fourth, agents derive utility from transfers to dependent agents in the economy, rather than, e.g., transfers to their own children, or their own children's consumption. This allows orphans to consume; recall that agents face mortality risk through their whole lives, so some dependent children will have no living parents.

Finally, some remaining assumptions, e.g. log utility, are quite standard.

One can discuss how plausible each of these assumptions is, but together they do generate the fertility behavior described by (10), and may at any rate serve as a good starting point.

Consider thus the following set-up. In each model period t, agents in life periods $j \in \{\underline{B}, ..., \overline{B}\}$ allocate the wage $w_{j,t}$ between own consumption in the same period, $c_{j,t}$, total transfers to non-active agents, $\tau_{j,t}$, and spending on conceiving $n_{j,t}$ children, at total cost $q_j n_{j,t}^{1/\delta}$. As in Section 3.1, δ is the inverse of the elasticity of the child-cost function, and q_j is a parameter that shifts the cost function and is here age dependent. This gives the budget constraint

$$c_{j,t} = w_{j,t} - \tau_{j,t} - q_j n_{j,t}^{\frac{1}{\delta}}. \tag{A1}$$

For simplicity, utility is defined over the economically and reproductively active phase of life only; all other agents consume their incomes or transfers received, thus making no choices. In particular, agents who work but have no children ($j \in \{\overline{B}+1, ..., R\}$) consume all their income and do not make transfers. This can easily be relaxed, by defining their utility in terms of transfers and consumption only, but not fertility.

An agent who becomes active in period t has utility

$$U_t = \sum_{j=\underline{B}}^{\overline{B}} \beta^{j-\underline{B}} u_{j,t+(j-\underline{B})}, \tag{A2}$$

where $u_{j,t}$ is direct utility (j being the period of life, t the model period), and $\beta \in (0,1)$ is a discount factor. The agents care about their own consumption, transfers, and the number of children conceived, with weights η and $\widetilde{\gamma}$ on the latter two. The period-utility can thus be written

$$u_{j,t} = (1 - \eta - \widetilde{\gamma}) \ln(c_{j,t}) + \eta \ln(\tau_{j,t}) + \widetilde{\gamma} \ln(n_{j,t}). \tag{A3}$$

Because agents have no means of allocating resources intertemporally, they simply maximize the direct utility in every period, as given by (A3), subject to the budget constraint in (A1). The first-order conditions with respect to $n_{j,t}$ and $\tau_{j,t}$ state that

$$(1 - \eta - \widetilde{\gamma}) \left[w_{j,t} - \tau_{j,t} - q_j n_{j,t}^{\frac{1}{\delta}} \right]^{-1} \left(\frac{1}{\delta} \right) q_j n_{j,t}^{\frac{1}{\delta}-1} = \widetilde{\gamma} \left[n_{j,t} \right]^{-1}, \tag{A4}$$

and

$$(1 - \eta - \widetilde{\gamma}) \left[w_{j,t} - \tau_{j,t} - q_j n_{j,t}^{\frac{1}{\delta}} \right]^{-1} = \eta \left[\tau_{j,t} \right]^{-1}. \tag{A5}$$

These can be solved to give consumption as a constant fraction of income:

$$c_{j,t} = \left[\frac{1 - \eta - \widetilde{\gamma}}{1 - \widetilde{\gamma}(1 - \delta)} \right] w_{j,t}. \tag{A6}$$

Using (A1), (A4), and (A6), spending on children also becomes a constant fraction of income, $q_j n_{j,t}^{1/\delta} = \{ \delta \widetilde{\gamma} / [1 - \widetilde{\gamma}(1 - \delta)] \} w_{j,t}$, so optimal fertility can be written as in (10), where

$$\gamma_j = \left(\frac{\delta \widetilde{\gamma}}{q_j \left[1 - \widetilde{\gamma}(1 - \delta) \right]} \right)^{\delta}. \tag{A7}$$

for $j \in \{\underline{B}, ..., \overline{B}\}$ and otherwise $\gamma_j = 0$.

Moreover, using (A4) and (A5), transfers become a constant fraction of income, $\tau_{j,t} = \{ \eta / [1 - \widetilde{\gamma}(1 - \delta)] \} w_{j,t}$, so that the factor $\eta / [1 - \widetilde{\gamma}(1 - \delta)]$ effectively functions as a voluntary tax rate.

For internal consistency of the model, we may also specify consumption of economically non-active agents. Dependent children simply rely on transfers from the active generations. Because the recipients of these transfers do not reproduce, the transfers play no role for the results. To model consumption of retired agents we can make several assumptions. For example, they could be partial recipients of the same transfers that children receive. Here, we let retired agents be owners of land from which they earn rental income. This interpretation means that agents receive an amount of land when they stop working as inheritance from landowners who die. (This includes all agents in period T of life, but also some younger landowning agents, since agents die throughout the life cycle.) Like dependent children, old agents are also reproductively non-active so their incomes and consumption play no role for the results.

More compactly, this can be written as:

$$c_{j,t} = \begin{cases} c_{j,t}^{\text{dependent}} & \text{if } j \in \{1, ..., \underline{B} - 1\}, \\ c_{j,t}^{\text{reproductive}} & \text{if } j \in \{\underline{B}, ..., \overline{B}\}, \\ w_{j,t} & \text{if } j \in \{\overline{B} + 1, ..., R\}, \\ c_{j,t}^{\text{retired}} & \text{if } j \in \{R + 1, ..., T\}, \end{cases} \tag{A8}$$

where the notation is self-explanatory:

$$c_{j,t}^{\text{dependent}} = \frac{\sum_{j=\underline{B}}^{\overline{B}} \tau_{j,t} P_{j,t}}{\sum_{j=1}^{\underline{B}-1} P_{j,t}} = \left[\frac{\eta}{1 - \widetilde{\gamma}(1 - \delta)}\right] \frac{\sum_{j=\underline{B}}^{\overline{B}} w_{j,t} P_{j,t}}{\sum_{j=1}^{\underline{B}-1} P_{j,t}}, \tag{A9}$$

$$c_{j,t}^{\text{reproductive}} = \left[\frac{1 - \eta - \widetilde{\gamma}}{1 - \widetilde{\gamma}(1 - \delta)}\right] w_{j,t}, \tag{A10}$$

$$c_{j,t}^{\text{retired}} = \frac{Y_t - \sum_{j=\underline{B}}^{R} w_{j,t} P_{j,t}}{\sum_{j=R+1}^{T} P_{j,t}} = \frac{\alpha Y_t}{\sum_{j=R+1}^{T} P_{j,t}}. \tag{A11}$$

The last of these uses (16).

B Data and calibration

B.1 Data on wages, fertility, and age-specific mortality

The age-specific fertility data in Figure 2 are from Statistics Sweden (1969, Table 34), and report the number of children born by women in seven different five-year age brackets, the youngest 15-19 years old, and the oldest 45-49. These correspond to the periods (years) of life 16-20 and 46-50, respectively. The data refer to the mother's age when giving birth, while the j in $n_{j,t}$ in the model refers to the period of life in which the child is conceived. Because a child is born in the period after it is conceived, the model and the data coincide numerically: for example, children conceived in period of life $j = 15$ are born when the mother is of age 15.

The age-specific wages in Figure 2 are from Statistics Sweden (1942, Table 11). These are annual wages in 1940 for male rural workers (*lantarbetare*), who were paid some of their wages in kind (*arbetare i kost*), and the wage includes the in-kind payment. 1940 is the earliest year available in these data, and coincides with the Second World War, a period when Sweden was largely cut off from trade, which may partly help mimic preindustrial conditions.

The wage rates are defined for slightly different age brackets than for fertility. The youngest and oldest brackets are open ended, labelled 18 and below, and 65 and above, respectively. We assume that the youngest and oldest workers are of ages 14 and 69, respectively, corresponding to 15 and 70 in terms of periods of life; recall that we set $\underline{B} = 15$ and $R = 70$. The wage of the youngest working age group (in life periods 15-19) is normalized to one.

The age-specific survival rates in Table 2 are from Statistics Sweden (1969, Table 40) and refer to the Swedish total population 1750-1800. The rates are defined for different age brackets, the youngest of age 0 (the first year of life), the second youngest of ages 1-2 (second and third years of life), and so on. The last age bracket is 80 and above, which is here translated to ages 80-89, meaning that the 90th period of life is the last; recall that we set $T = 90$.

B.2 Setting γ_j

Let w_i^{data} be the wage in the data referring to an agent in the ith period of life (of age $i-1$), and let n_i^{data} be the corresponding number from the fertility data (the number of children conceived in the ith period of life). The age-dependent fertility parameters are first set such that, for $j \in \{15, 19\}$, γ_j is given by

$$\gamma_j = \frac{\sum_{i=15}^{19} n_i^{\text{data}}}{\sum_{i=15}^{19} \left(w_i^{\text{data}}\right)^{\delta}}, \tag{A12}$$

and analogously for the other age groups. The series of γ_j's can be scaled uniformly to make the age-wage profile match the data in Figure 2. As can be understood intuitively from (9) in the simplified model, these parameters are inversely related to steady-state wages. In the benchmark specification no scaling is used. Values of γ_j for select age groups are shown in Table 3.

B.3 Setting β_j and ρ

We set $\rho = .9$, close to perfect substitutability.

We let β_j increase by 1% per year of life from $\underline{B} = 15$ until period 32 (age 31), after which β_j decreases by 1% per year of life and drops to zero after period $R = 70$, all subject to $\sum_{j=\underline{B}}^{R} \beta_j = 1$. Values for select age groups are shown in Table 3.

The period of life when β_j peaks is chosen to roughly coincide with the peak in the age-wage profile in the data. Note that wages decline more slowly with age than they rise. This reflects that older cohorts are smaller, due to mortality throughout the life cycle, making older agents scarcer than young, and thus their marginal products higher; this is where imperfect substitutability ($\rho < 1$) plays a role.

B.4 Elasticities

Setting $\delta = .15$ translates directly to an elasticity of *conceived* children with respect to same-year wages of .15; see (10). To compare this to data, we compute the elasticity of the Crude Birth Rate in a given year—the number of births over population—with respect to wages in the previous year. The CBR in year t equals the number newborns (agents who are in the first period of life), over the total population in the same period, excluding the newborns themselves:

$$\text{CBR in year } t = \frac{P_{1,t}}{\sum_{i=2}^{T} P_{i,t}}. \tag{A13}$$

The Crude Death Rate is computed as the total number of agents who die in a given year, divided by total population:

$$\text{CDR in year } t = \frac{\sum_{i=1}^{T}(1 - s_{i.t})P_{i,t}}{\sum_{i=1}^{T} P_{i,t}}. \tag{A14}$$

Wages in year t are computed as the (unweighted) average across all age groups of working age:

$$\text{Wage in year } t = \frac{\sum_{i=\underline{B}}^{R} w_{i,t}}{R + 1 - \underline{B}}. \tag{A15}$$

We then simulate the model, under the benchmark parametrization, for 501 years (corresponding to 1300-1800), and 100 times. The result is a panel data set with variation in wages, CBR, and CDR across different (artificial) economies and over time. We then run a number of fixed-effects regressions with the logarithms of CBR and CDR as dependent variables, and log wages at different lags as independent variables.

The results are shown in Table A.1 for log wages at different lags. The reported coefficient estimates represent elasticities since all variables are logged. The table also reports the sum of the different coefficients, similar to the empirical literature (e.g., Klemp and Møller 2016, Lagerlöf 2015). Regardless of specification, the estimates are around .13 and $-.2$ for CBR and CDR, respectively.

Note that we have not imposed anything on the model to produce the measured effects from lagged variables. Rather, these effects are driven by the serially correlated productivity shocks and the way the shocks are propagated through the model.

The CDR elasticity is guided by κ, which was set to .01, and measures the elasticity of the survival rate from starvation, s_t^y, with respect downward changes in per-capita incomes; see (22). The estimated CDR elasticity is much larger in absolute terms than the survival elasticity, because the survival rates are close to one: small percentage changes in survival rates correspond to large percentage changes in mortality rates.

C Initial conditions

We set initial output and population levels of each age group to those associated with the steady state to which the economy converges when productivity and mortality are constant over time, i.e., $\sigma_A = \sigma_X = \phi = \kappa = \mu_t = 0$, which implies $A_t = X_t = 1$, and $s_{j,t} = s_j$, for all t.

To find analytical solutions requires fertility to be linear in wages ($\delta = 1$). Because of this, and because we have closed down the shocks, the steady state distribution of the different variables will change over the first several periods. We therefore disregard the first 500 periods when calculating the moments.

First note that, with $\delta = 1$, fertility in (10) can be written

$$n_{j,t} = \gamma_j w_{j,t} = \gamma_j (1-\alpha) \frac{Y_t}{L_t} \beta_j \left(\frac{L_t}{P_{j,t}} \right)^{1-\rho}, \tag{A16}$$

where the second equality uses the expression for wages in (15). Since $\gamma_j = n_{j,t} = 0$ for all $j \notin \{\underline{B}, ..., \overline{B}\}$, we can write the sum in (17) as

$$P_{1,t+1} = \sum_{j=\underline{B}}^{\overline{B}} n_{j,t} P_{j,t} = \sum_{j=1}^{T} n_{j,t} P_{j,t}. \tag{A17}$$

Similarly, since $\beta_j = 0$ for all $j \notin \{\underline{B}, ..., R\}$, we can write the sum in (14) as

$$L_t = \left[\sum_{j=\underline{B}}^{R} \beta_j P_{j,t}^{\rho} \right]^{\frac{1}{\rho}} = \left[\sum_{j=1}^{T} \beta_j P_{j,t}^{\rho} \right]^{\frac{1}{\rho}}. \tag{A18}$$

Now, (A16) to (A18) give the new-born population in period $t+1$ as

$$\begin{aligned} P_{1,t+1} &= \textstyle\sum_{j=1}^{T} \gamma_j (1-\alpha) \frac{Y_t}{L_t} \beta_j \left(\frac{L_t}{P_{j,t}} \right)^{1-\rho} P_{j,t} \\ &= (1-\alpha) Y_t \frac{\sum_{j=1}^{T} \gamma_j \beta_j P_{j,t}^{\rho}}{L_t^{\rho}} \\ &= (1-\alpha) \theta_t Y_t, \end{aligned} \tag{A19}$$

where

$$\theta_t = \frac{\sum_{j=1}^{T} \gamma_j \beta_j P_{j,t}^{\rho}}{\sum_{j=1}^{T} \beta_j P_{j,t}^{\rho}}. \tag{A20}$$

Setting $s_{j,t} = s_j$ for all t, population dynamics in (18) can be written $P_{j+1,t+1} = s_j P_{j,t}$. This difference equation can be solved to give

$$P_{j,t} = S_j P_{1,t-j+1}, \tag{A21}$$

where

$$S_j = \prod_{i=1}^{j-1} s_i \tag{A22}$$

is the (time invariant) probability that an agent survives to the jth period of life, and where $S_1 = 1$ and $P_{1,0}$ are given. Using (A20) and (A21) gives

$$\theta_t = \frac{\sum_{j=1}^{T} \gamma_j \beta_j P_{j,t}^{\rho}}{\sum_{j=1}^{T} \beta_j P_{j,t}^{\rho}} = \frac{\sum_{j=1}^{T} \gamma_j \beta_j S_j^{\rho} P_{1,t-j+1}^{\rho}}{\sum_{j=1}^{T} \beta_j S_j^{\rho} P_{1,t-j+1}^{\rho}}. \tag{A23}$$

In a steady-state equilibrium, where $P_{1,t}$ is constant at $\overline{P}_1$, (A23) can be written

$$\overline{\theta} = \frac{\overline{P}_1^{\rho} \sum_{j=1}^{T} \gamma_j \beta_j S_j^{\rho}}{\overline{P}_1^{\rho} \sum_{j=1}^{T} \beta_j S_j^{\rho}} = \frac{\sum_{j=1}^{T} \gamma_j \beta_j S_j^{\rho}}{\sum_{j=1}^{T} \beta_j S_j^{\rho}}. \tag{A24}$$

Imposing steady state on (A21) gives

$$\overline{P}_j = S_j \overline{P}_1, \tag{A25}$$

so the steady-state level of effective labor in (14) becomes:

$$\overline{L} = \left[\sum_{j=1}^{T} \beta_j \overline{P}_j^{\rho} \right]^{\frac{1}{\rho}} = \left[\sum_{j=1}^{T} \beta_j S_j^{\rho} \overline{P}_1^{\rho} \right]^{\frac{1}{\rho}} = \overline{P}_1 \left[\sum_{j=1}^{T} \beta_j S_j^{\rho} \right]^{\frac{1}{\rho}}. \tag{A26}$$

Using (A26), and setting $A_t = X_t = 1$ in (11), gives steady-state output as

$$\overline{Y} = \overline{L}^{1-\alpha} = \overline{P}_1^{1-\alpha} \left[\sum_{j=1}^{T} \beta_j S_j^{\rho} \right]^{\frac{1-\alpha}{\rho}}. \tag{A27}$$

Evaluating $P_{1,t+1} = (1-\alpha)\theta_t Y_t$ from (A19) in steady state gives

$$\overline{P}_1 = (1-\alpha)\overline{\theta Y}. \tag{A28}$$

Now, (A27) and (A28) give

$$\overline{Y} = \left[(1-\alpha)\overline{\theta} \left(\sum_{j=1}^{T} \beta_j S_j^{\rho} \right)^{\frac{1}{\rho}} \right]^{\frac{1-\alpha}{\alpha}}, \tag{A29}$$

which together with (A22) and (A24) defines $\overline{Y}$ in terms of exogenous parameters, including the age-specific parameters γ_j, β_j, and s_j. Then the steady-state population of the new-born age group can be computed from (A28), and those of the other age groups from (A22) and (A25).

Using (A25) and (A29), steady-state output per agent can be written

$$\overline{y} = \frac{\overline{Y}}{\sum_{j=1}^{T} \overline{P}_j} = \frac{\overline{Y}}{\overline{P}_1 \sum_{j=1}^{T} S_j}. \tag{A30}$$

Tables

Parameter	**Value**	**Interpretation/comment**
T	90	Last period of life
$\underline{B}$	15	Period of life when agents become economically and reproductively active
$\overline{B}$	49	Last period of life being reproductively active
R	70	Last period of life being economically active
α	.4	Land share of output; same as Hansen and Prescott (2002)
δ	.15	Gives fertility-wage elasticity of .13; similar to Klemp and Møller (2016), Lagerlöf (2015)
κ	.01	Gives mortality-wage elasticity of $-.2$; similar to Klemp and Møller (2016), Lagerlöf (2015)
ϕ	.01	Expected death rate from disease is about 1% per year
γ_j	See Table 3	To match age-fertility profile to Swedish data in Figure 2
β_j	See Table 3	To match age-income profile to Swedish data in Figure 2
ρ	.9	Close to perfect substitutability between age cohorts when determining effective labor
μ_t	$.0125(t/500)^2$; $t = 0$ in 1300	Mean annual productivity growth rate rises at accelerating rate from 0 to 1.25 percent from 1300-1800; set to match population data in Figure 3
σ_A	.07	Standard deviation in permanent shock; set to match moments in the data
σ_X	.07	Standard deviation in temporary shock; set to match moments in the data

Table 1: Benchmark parameter values.

Ages $(j-1)$	**Survival rate in the data**	$E[s_t^d]s_j^{\text{age}}$
0	.797	.797
1-2	.947	.947
3-4	.972	.972
5-9	.987	.990
10-14	.993	.990
15-19	.993	.990
20-24	.992	.990
25-29	.990	.990
30-34	.988	.988
35-39	.988	.988
40-44	.984	.984
45-49	.982	.982
50-54	.978	.978
55-59	.973	.973
60-64	.959	.959
65-69	.941	.941
70-74	.908	.908
75-79	.870	.870
80-88	.775	.775
89	NA	0

Table 2: Age-specific survival probabilities, in Swedish data and in the model. The model values are the expected survival rate from age and disease, but disregarding starvation.

j	β_j	γ_j
18	0.0181	0.0104
32	0.0208	0.1128
40	0.0192	0.0468

Table 3: Values for γ_j and β_j for three select age groups.

Country	**Fraction years below 5th percentile**	**Fraction years above 95th percentile**
UK	.024	0
Netherlands	0	.099
Sweden	.017	0
Italy	0	.024
Portugal	.33	0
All	.054	.029

Table 4: Fraction of the years in which log GDP per capita fell below the 5th percentile, and above the 95th percentile, across 1000 simulated economies. The results are shown by country and for all country-years together.

Log GDP/capita time-series data

Country or robustness check	Mean	Standard deviation	1-year lag correlation	2-year lag correlation
England	−.16	.25	.96	.95
Netherlands	.31	.3	.96	.94
Italy	.18	.097	.85	.74
Sweden	.015	.13	.86	.74
Portugal	−.35	.17	.94	.92
Mean across five countries	.00	.19	0.91	.86

Simulations [5th percentile, 95th percentile]

Benchmark	[−.34, .45]	[.13, .42]	[.92, .99]	[.9, .99]
$\sigma_A = 0$, $\sigma_X = .35$	[−.062, .068]	[.18, .24]	[.37, .65]	[.36, .65]
Start year 1560	[−.29, .61]	[.094, .34]	[.87, .99]	[.82, .99]
$\delta = .3$	[−.25, .29]	[.13, .37]	[.93, .99]	[.9, .99]
$\kappa = .02$	[−.27, .37]	[.13, .39]	[.92, .99]	[.9, .99]
$\phi = .02$	[−.58, .63]	[.16, .55]	[.95, 1]	[.93, .99]

Table 5: Comparing moments in the data to the simulation results for the benchmark calibration and when changing some of the parameter values.

Figures

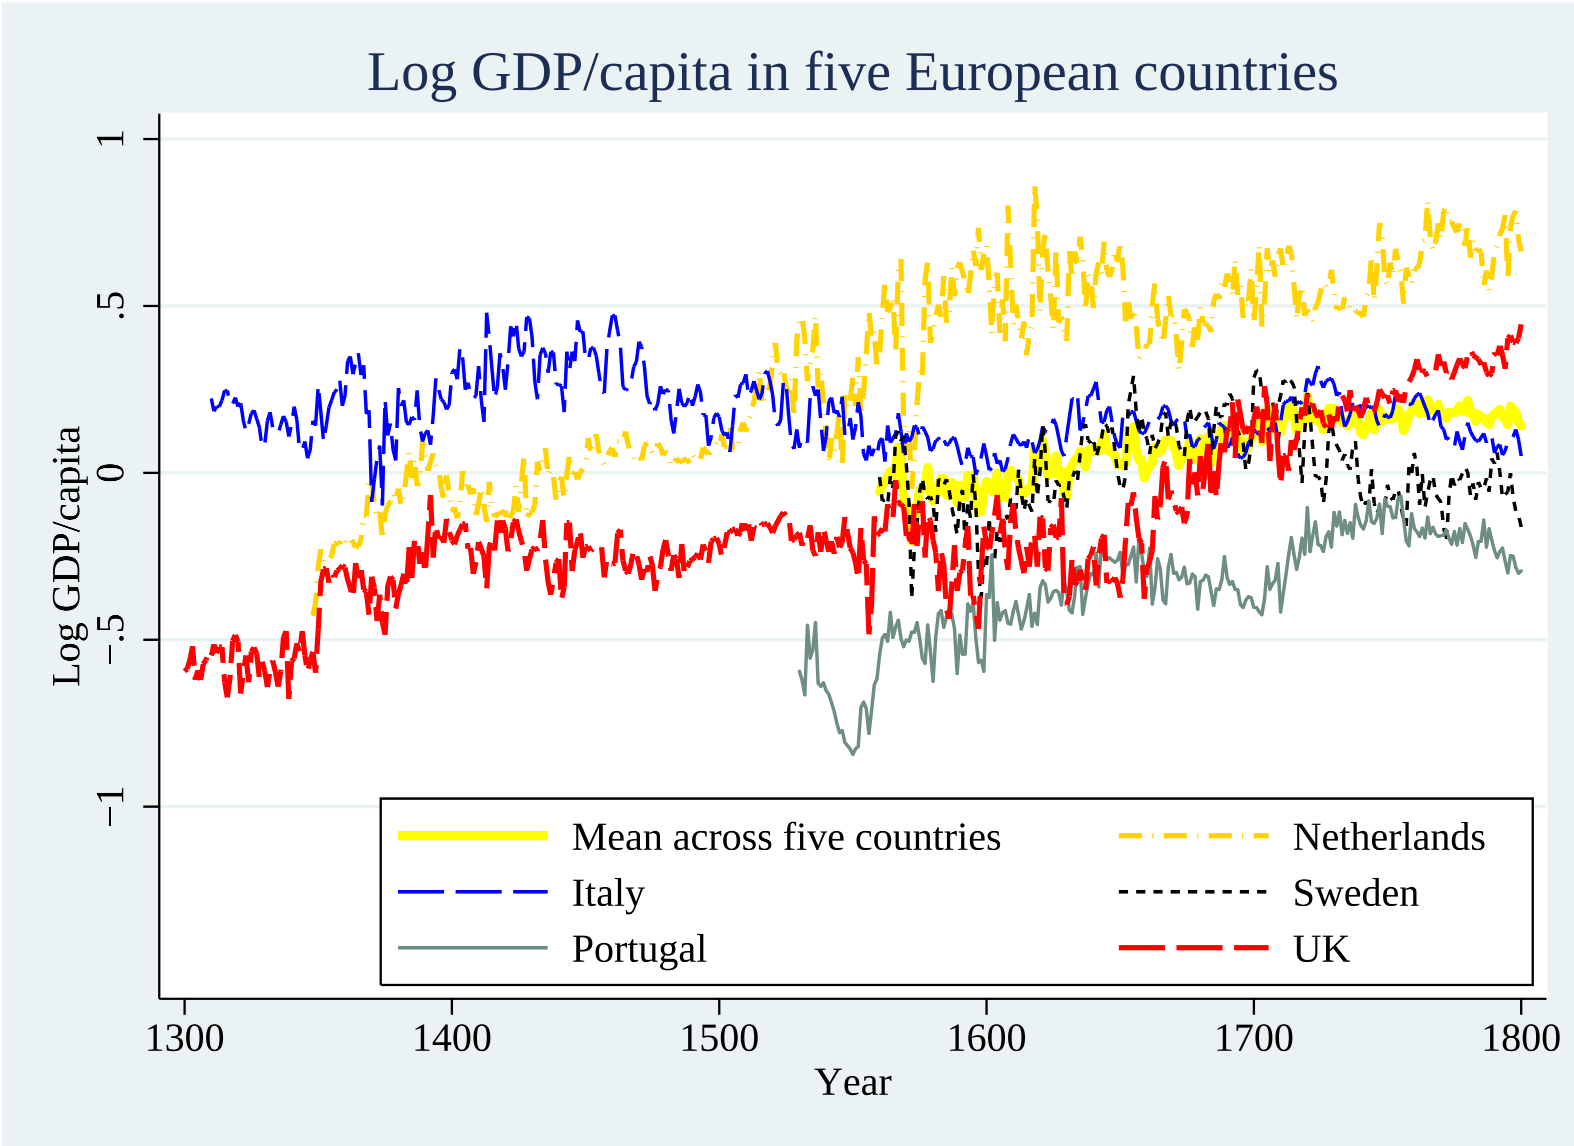

Figure 1: Data over log GDP per capita. The series are normalized to equal zero when averaged first over time for each country, and then across all countries.

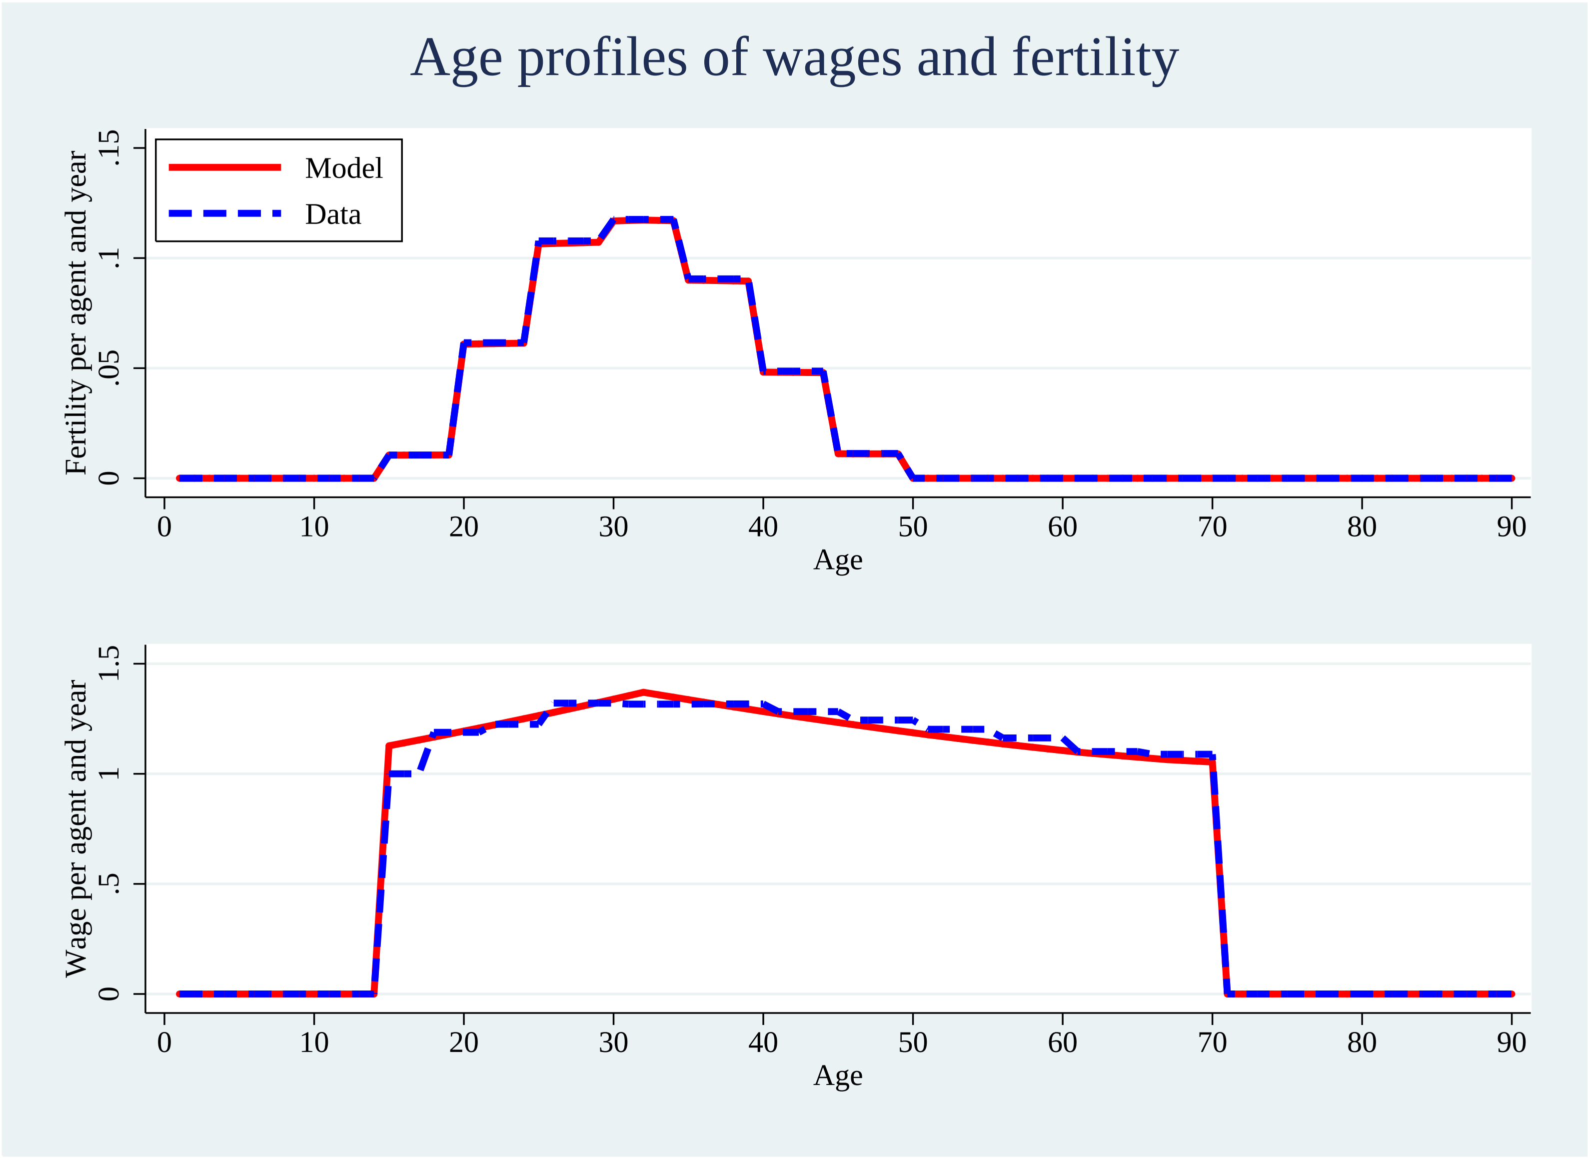

Figure 2: Fertility and wages by age group, in the data and in the model, guiding how γ_j, β_j and ρ are set. The data are from Sweden. Fertility rates are averages 1750-1800, and wages refer to rural Swedish workers in 1940, the earliest year available. The model-generated profiles are averages over 1000 runs for the last 50 years of the simulation, corresponding to the years 1750-1800.

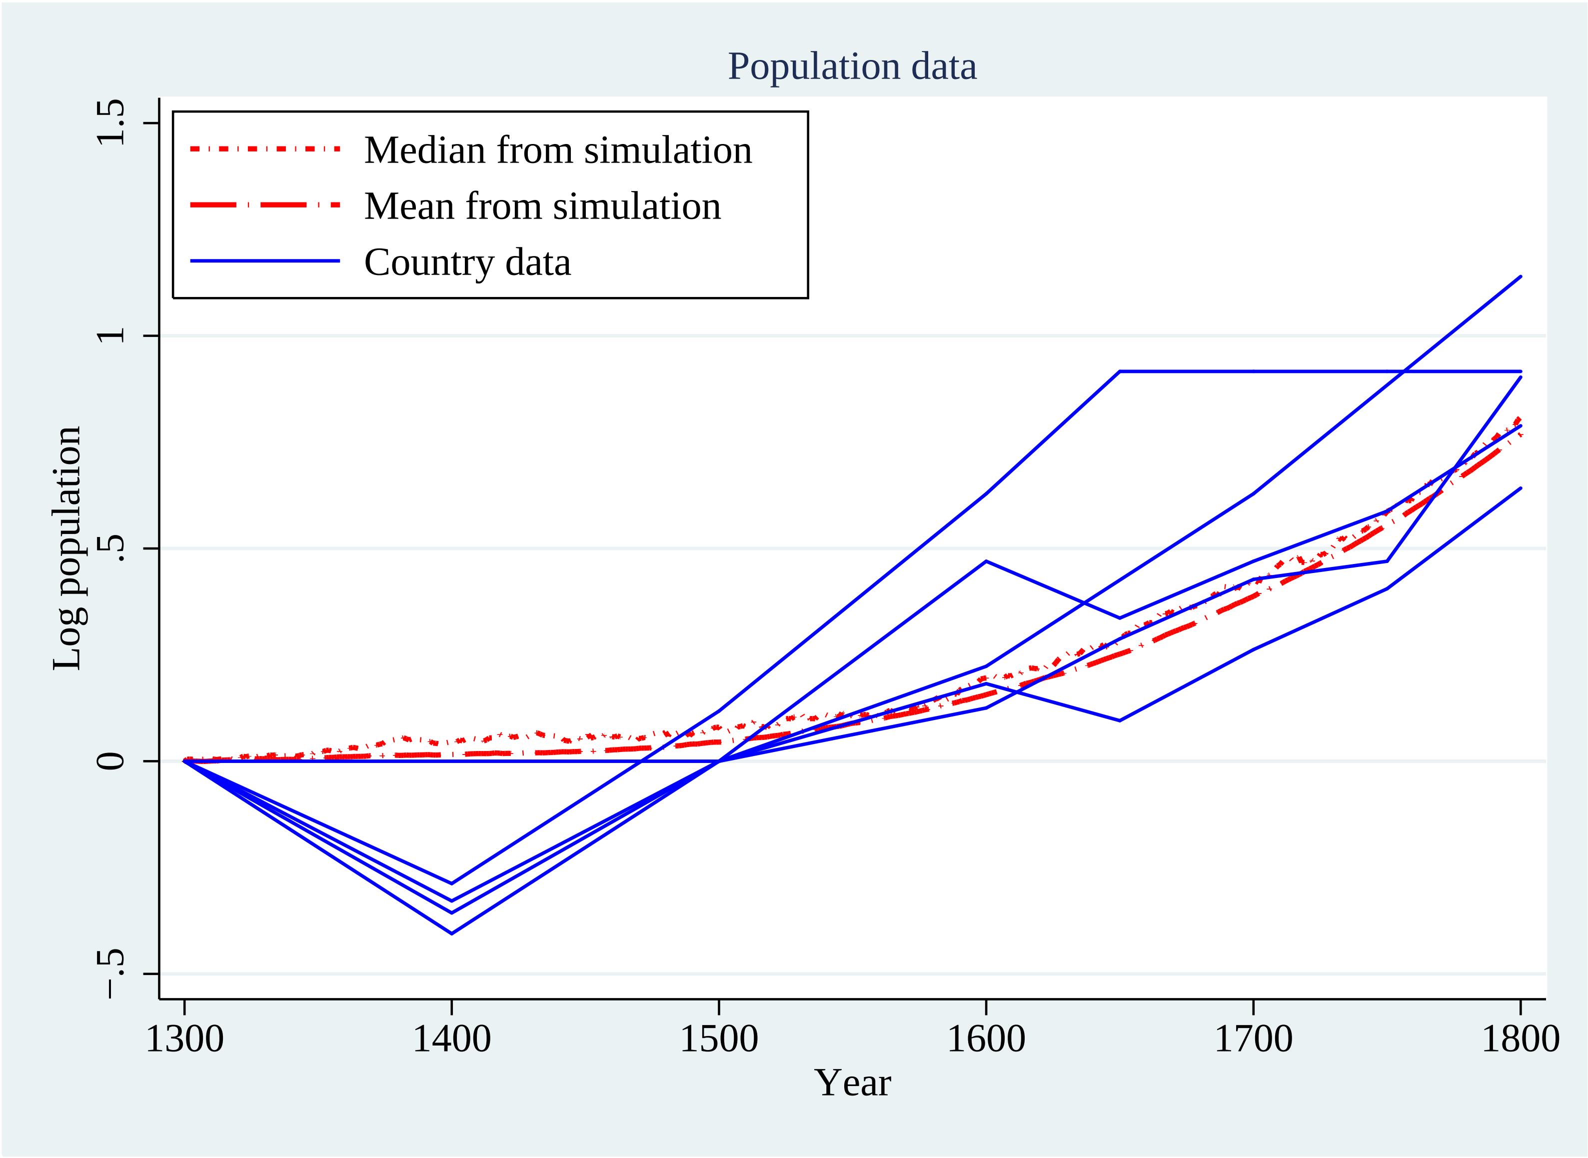

Figure 3: Population level data for the five countries from McEvedy and Jones (1978), and the mean and median population levels across 1000 simulated economies.

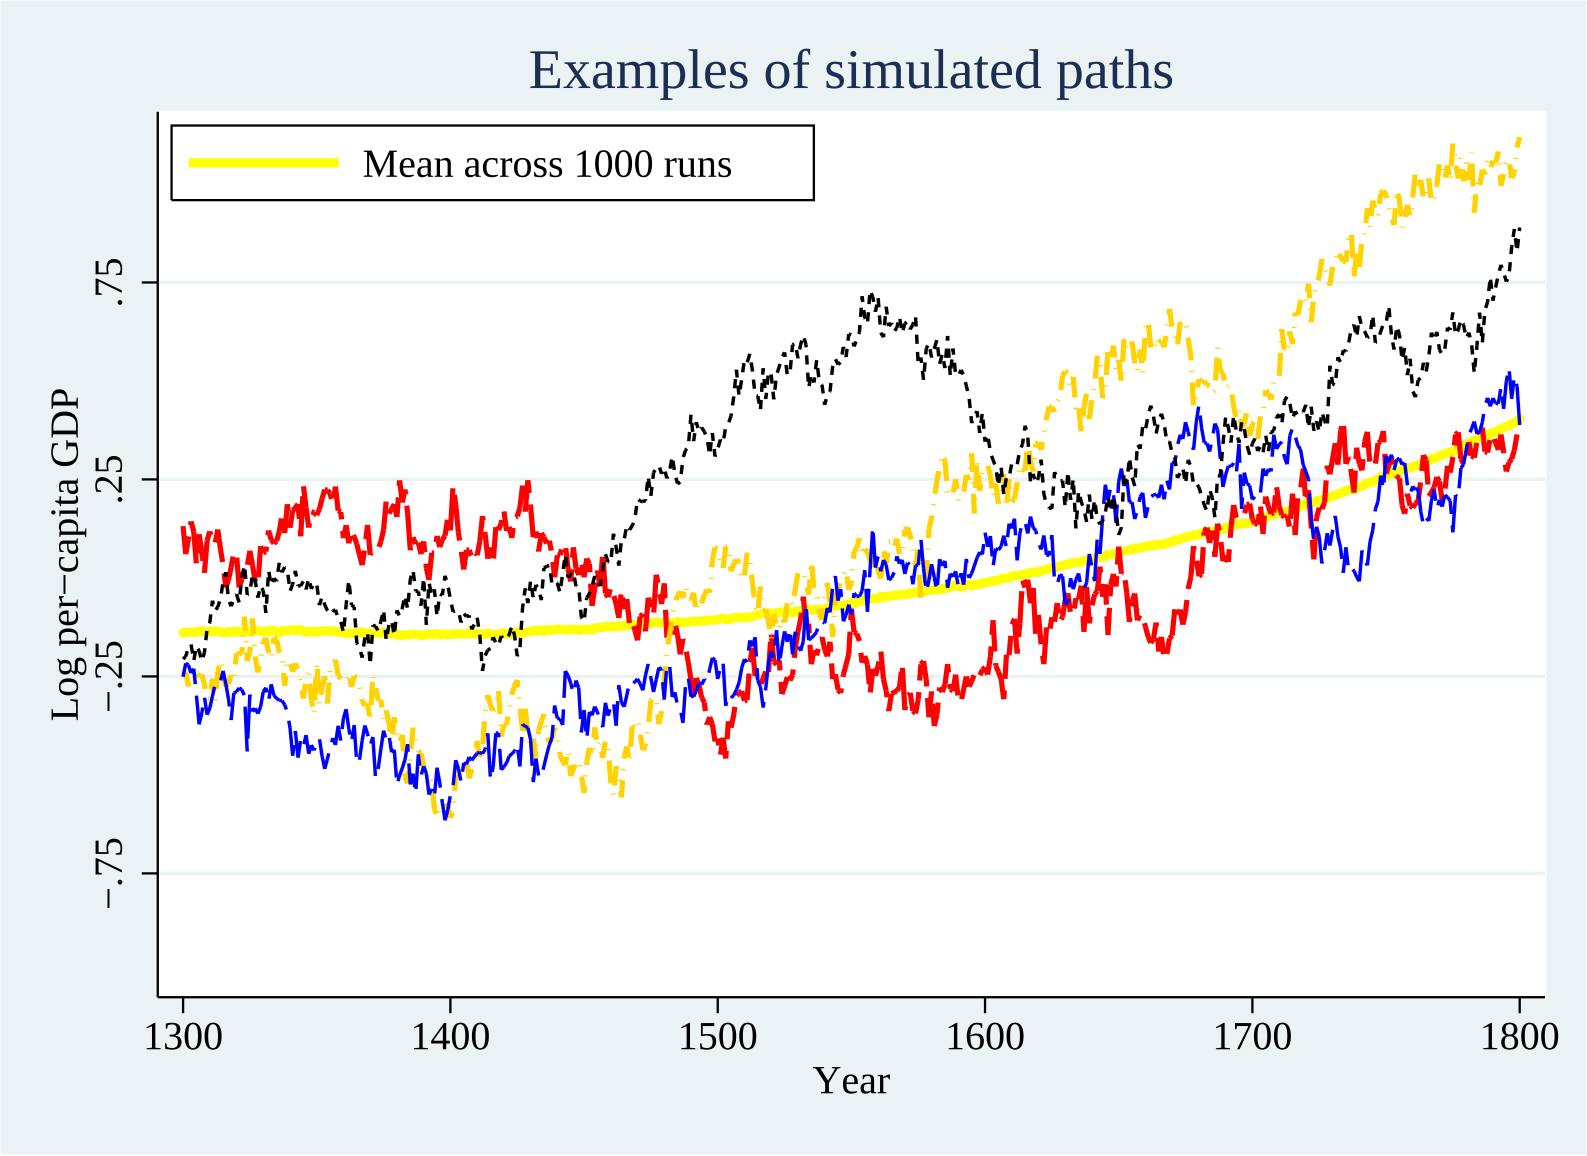

Figure 4: Examples of paths of log GDP per capita for four different simulated economies, and the average across 1000 runs (including the ones shown).

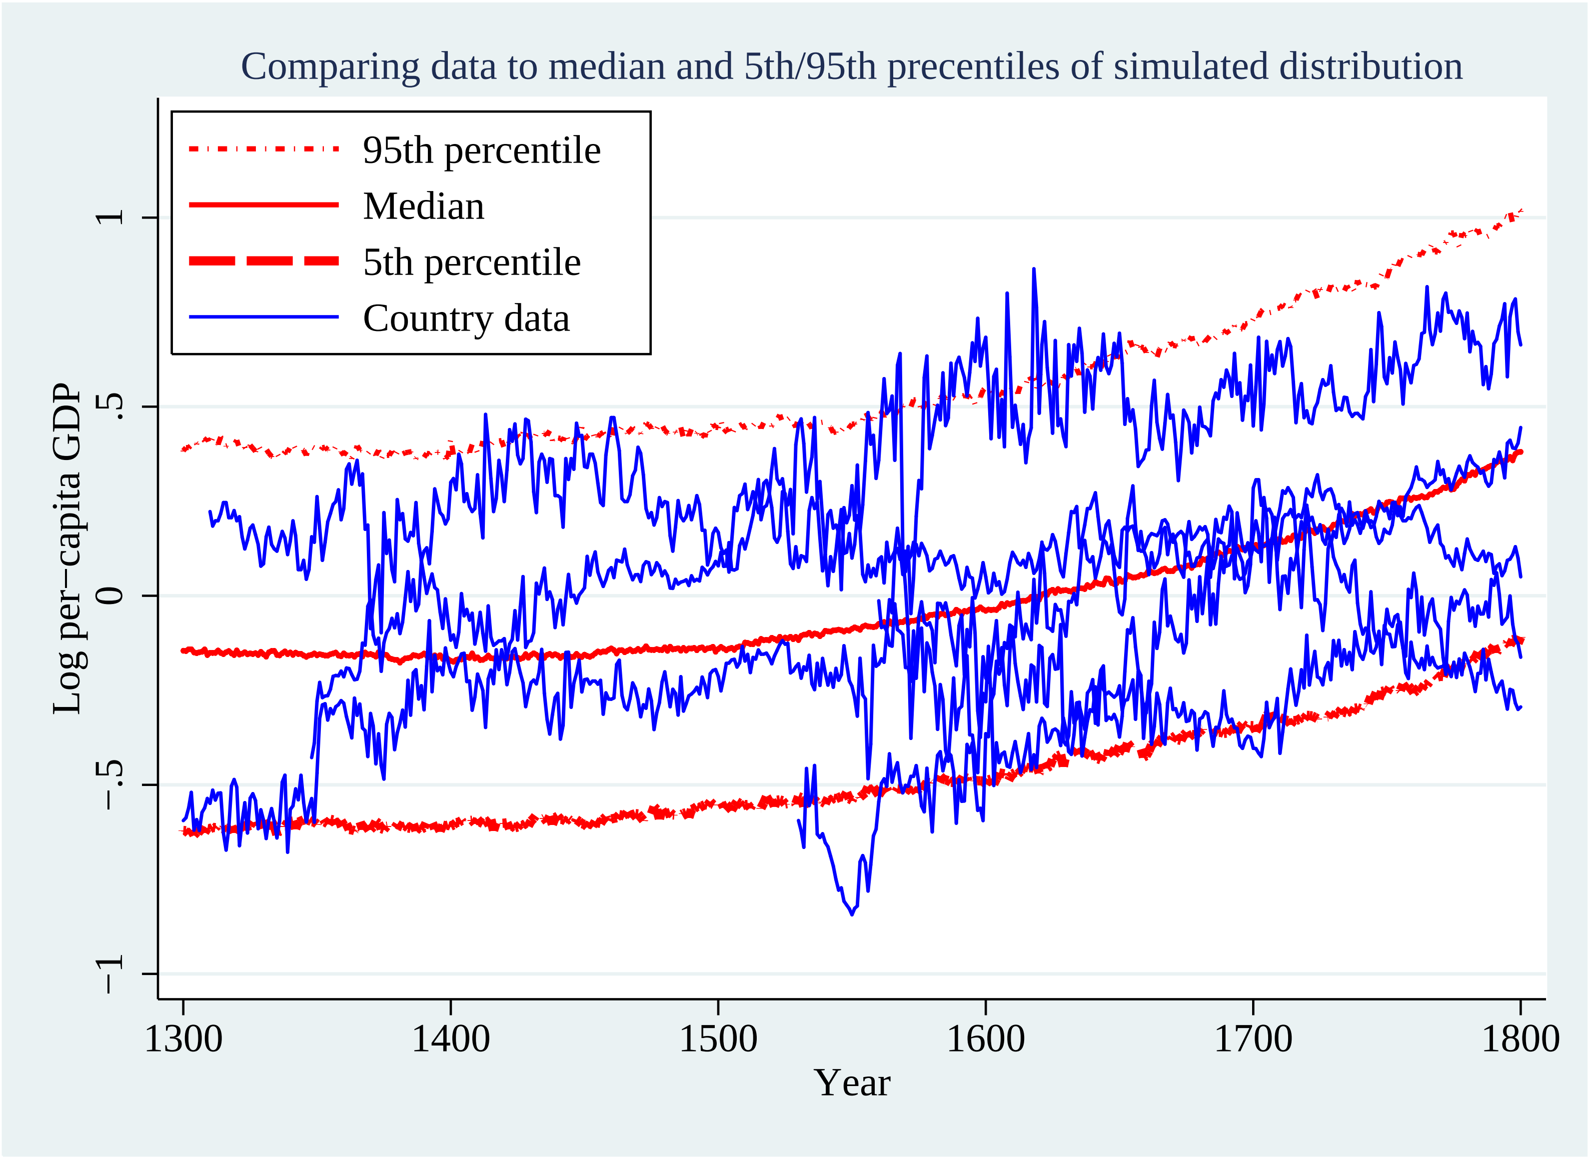

Figure 5: This figure compares the time paths in log GDP per capita for the five countries in Figure 1 to the median across 1000 runs by year, and the associated 5th and 95th percentiles for each year.

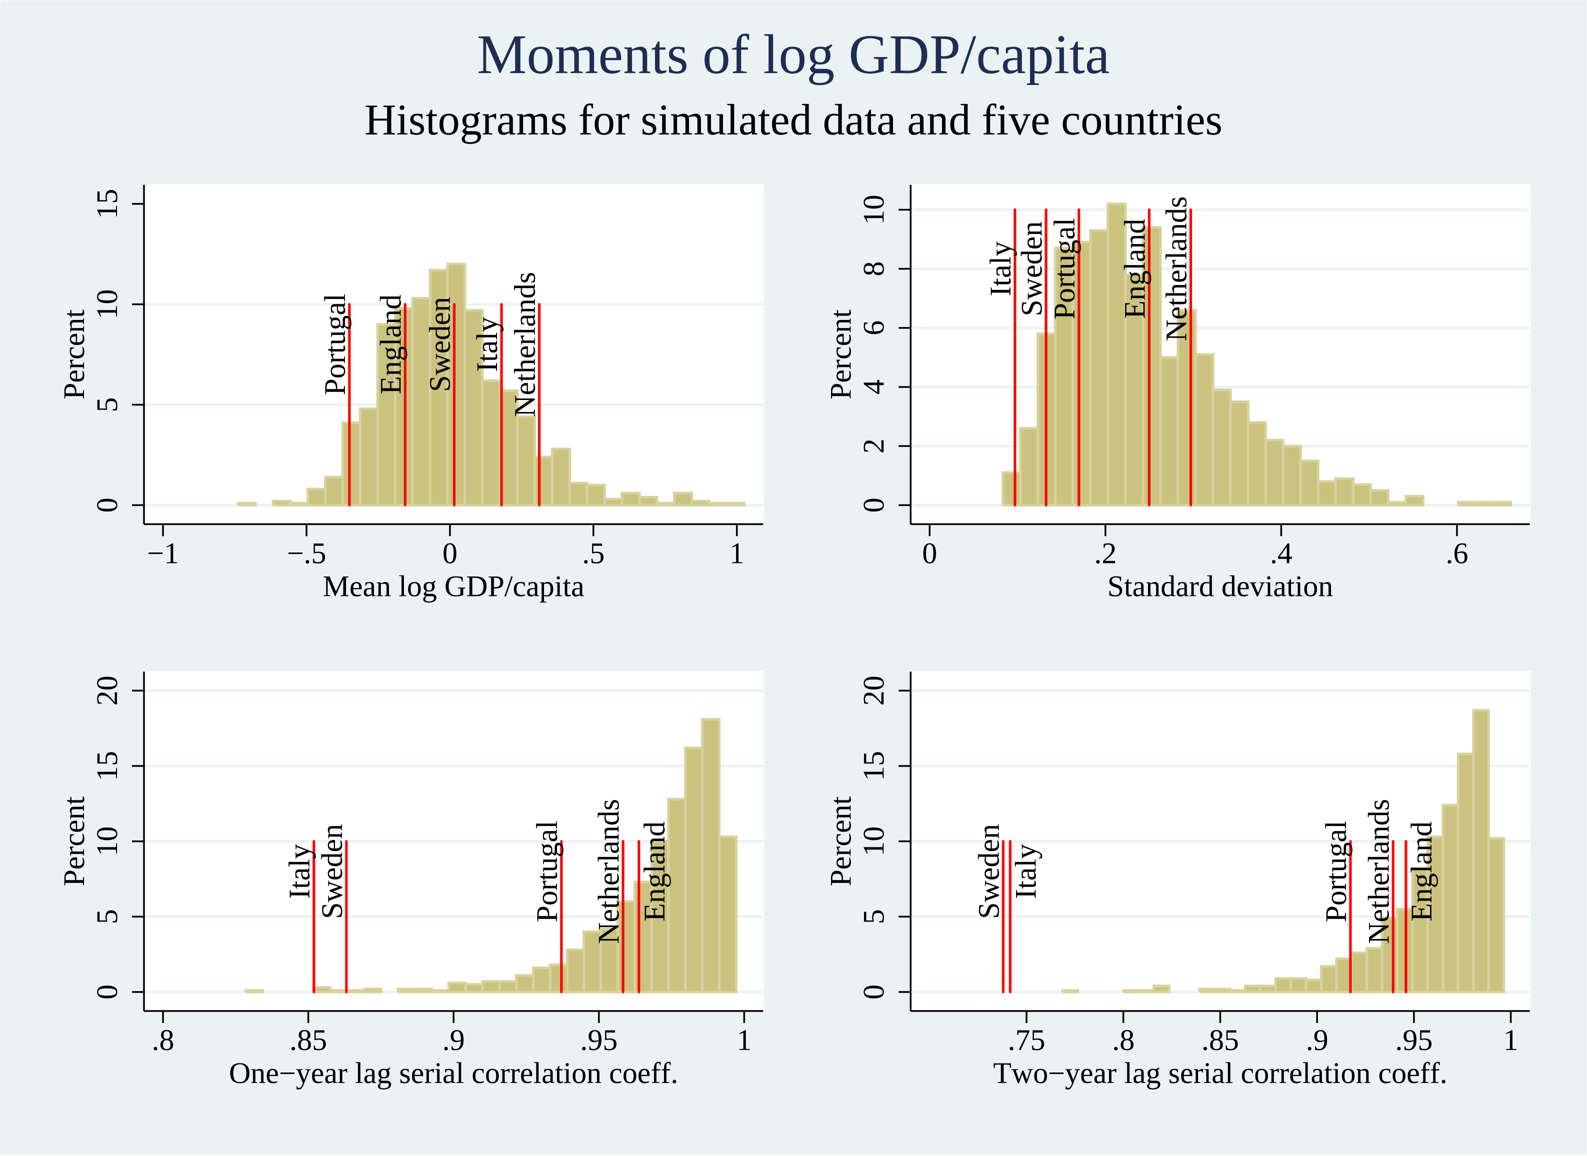

Figure 6: For each of the 1000 simulated economies, the means, standard deviations, and serial correlation coefficients (at one and two year lags) were calculated. This figure shows how these moments were distributed, as well as the corresponding values for the five countries.

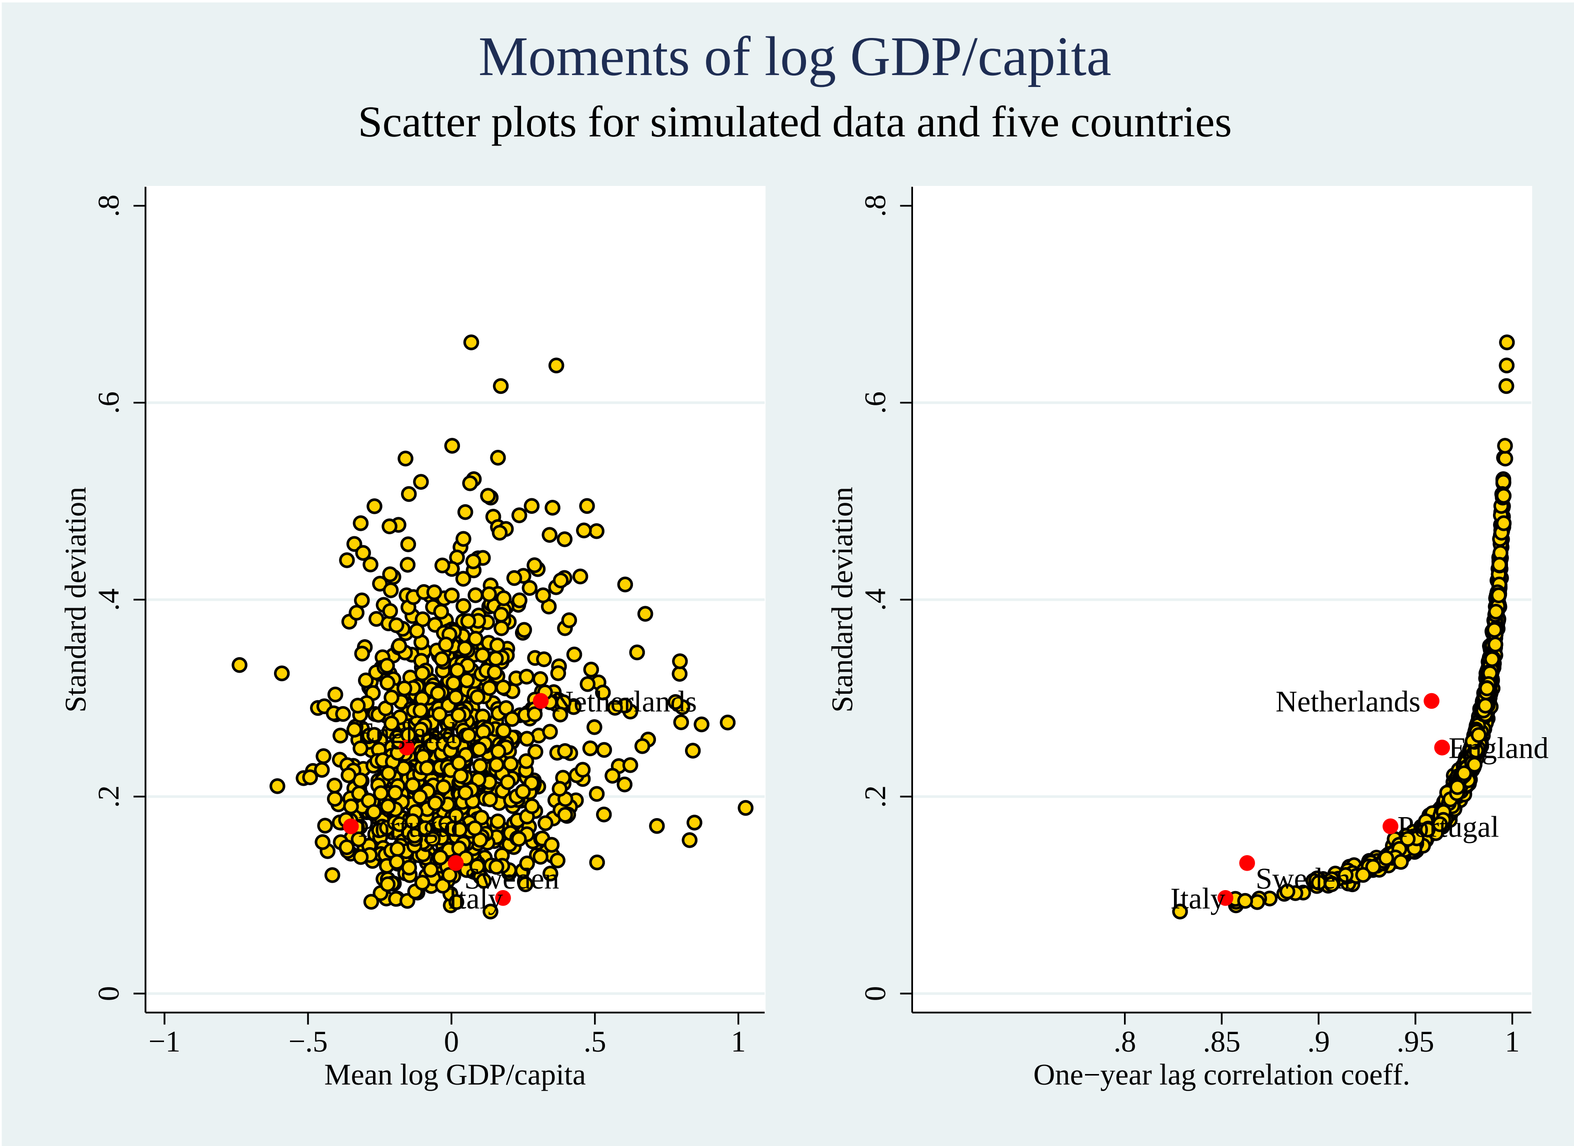

Figure 7: The same data as in Figure 6 presented in two plots.

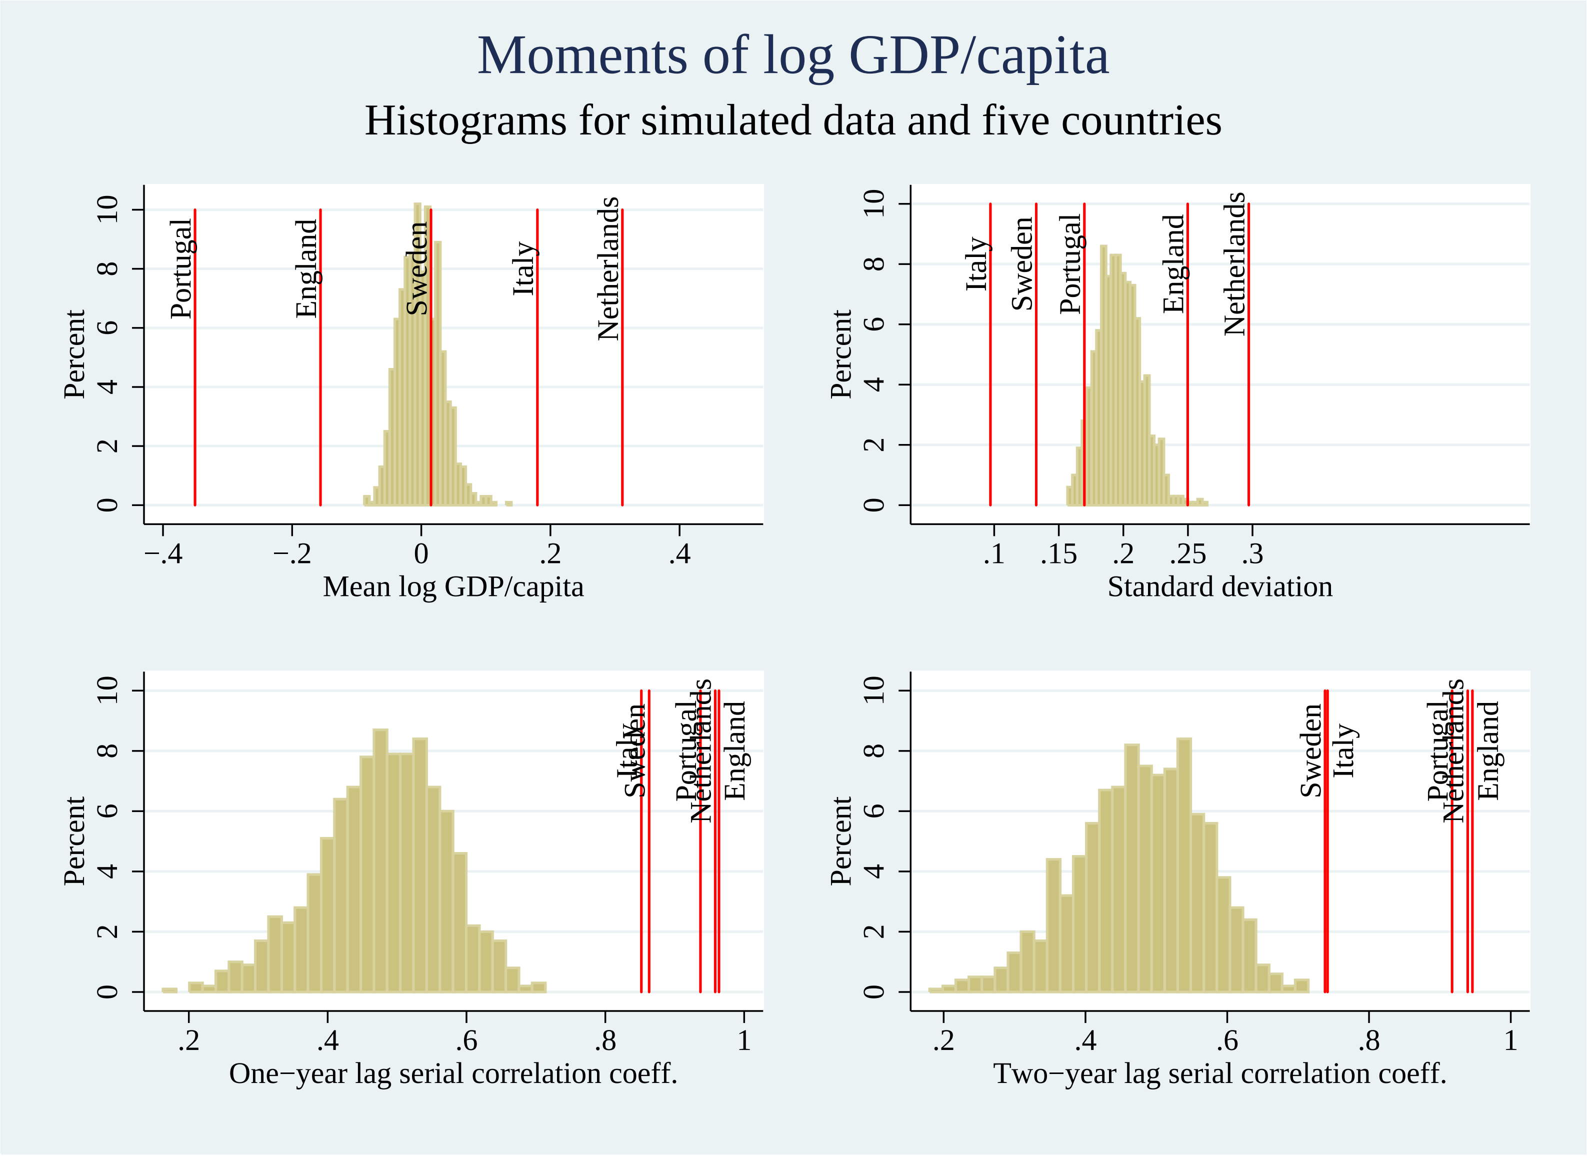

Figure 8: The same histograms as in Figure 6, when setting $\sigma_A = 0$ and $\sigma_X = .35$.

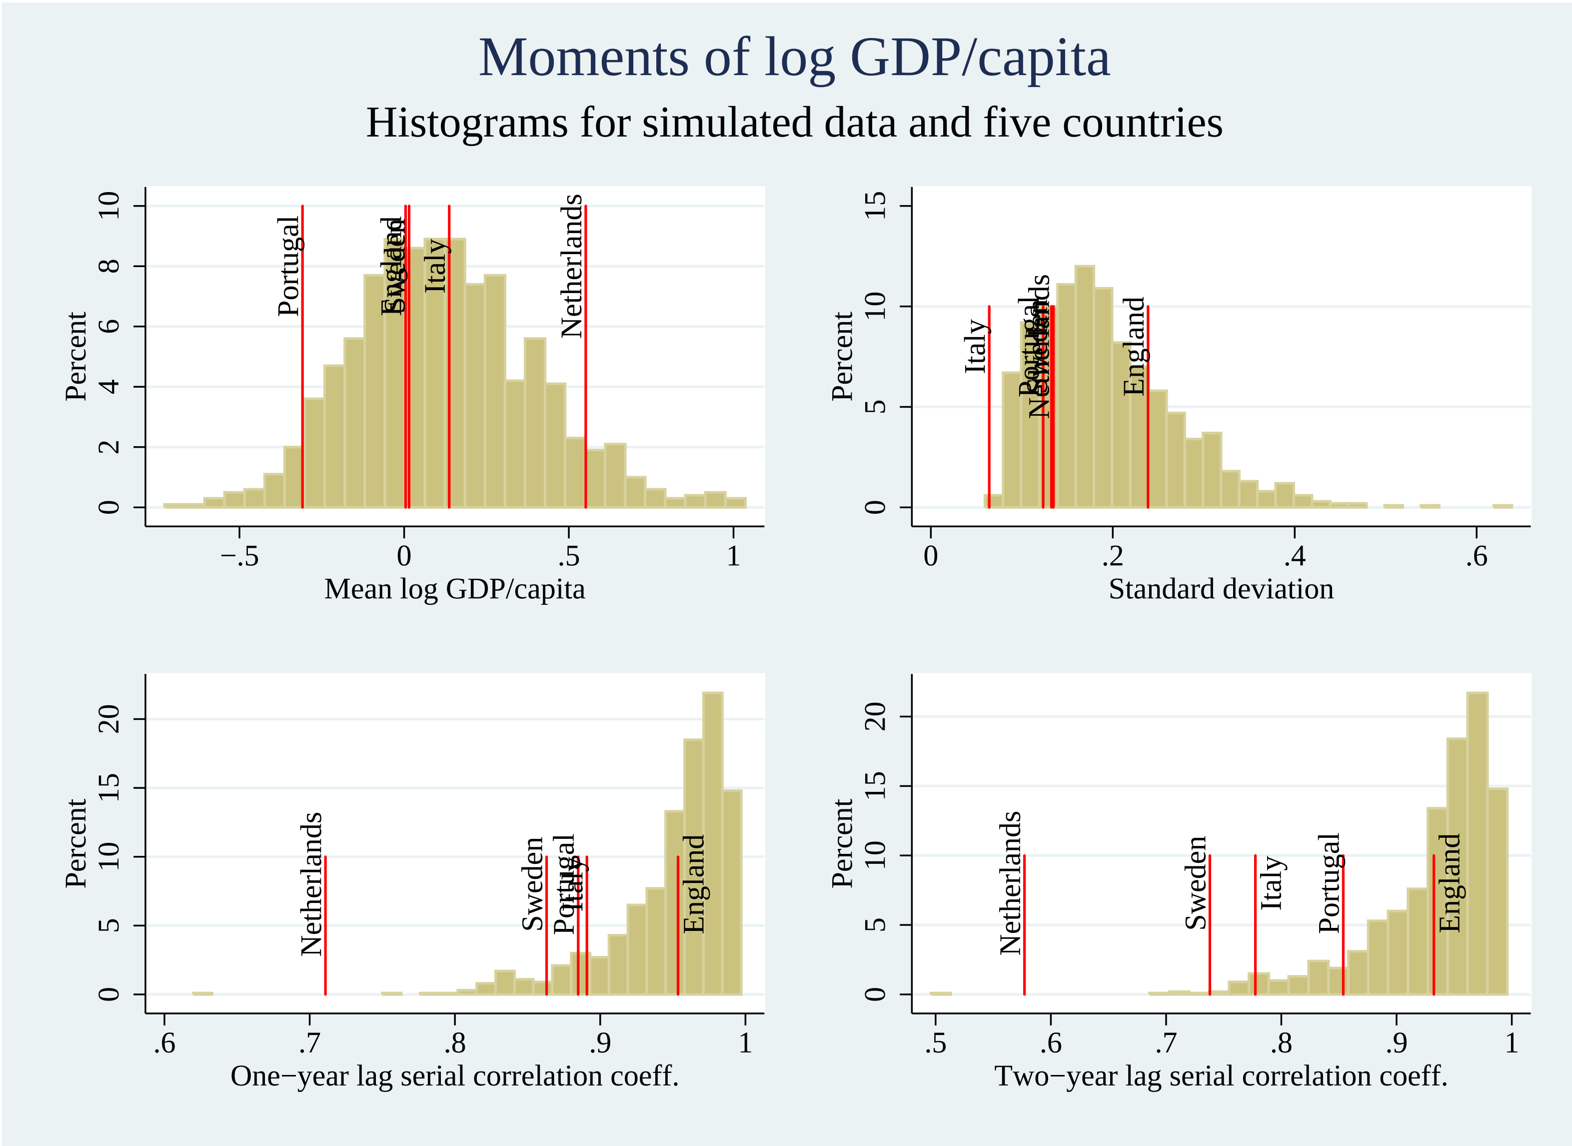

Figure 9: The same histograms as in Figure 6, considering the period 1560-1800, for which data are available for all five countries.

Appendix tables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Dep. variable is log CBR in year t				Dep. variable is log CDR in year t			
log wage t	0.130***			0.037***	−0.205***			−0.295***
	(382.96)			(25.26)	(−26.72)			(−8.46)
log wage $t-1$		0.132***		0.112***		−0.194***		0.113***
		(406.03)		(66.26)		(−25.21)		(2.81)
log wage $t-2$			0.128***	−0.017***			−0.195***	−0.021
			(360.06)	(−11.30)			(−25.34)	(−0.59)
Cumulative effect	0.130	0.132	0.128	0.132	−0.205	−0.194	−0.195	−0.203
Observations	50,100	50,000	49,900	49,900	50,100	50,000	49,900	49,900

Table A.1: Ordinary least squares regressions based on simulated data, with t-statistics in parentheses. The 50,100 observations consist of 100 economies, each simulated over 501 years (representing 1300-1800). The dependent variables are the crude birth and death rates, calculated as described in the Appendix. All regressions include fixed effects for year and simulated economy. * indicates $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.