Advances in Macroeconomics

Volume 1, Number 2 *2001* *Article 1*

Was an Industrial Revolution Inevitable? Economic Growth Over the Very Long Run

Charles I. Jones
University of California, Berkeley and NBER

Advances in Macroeconomics is one of *The B.E. Journals in Macroeconomics*, produced by bepress.com.

Copyright ©2001 by the authors.
All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher, bepress.com.

Was an Industrial Revolution Inevitable? Economic Growth Over the Very Long Run

Abstract

This paper studies a growth model that is able to match several key facts of economic history. For thousands of years, the average standard of living seems to have risen very little, despite increases in the level of technology and large increases in the level of the population. Then, after thousands of years of little change, the level of per capita consumption increased dramatically in less than two centuries. Quantitative analysis of the model highlights two factors central to understanding this history. The first is a virtuous circle: more people produce more ideas, which in turn makes additional population growth possible. The second is an improvement in institutions that promote innovation, such as property rights: the simulated economy indicates that arguably the single most important factor in the transition to modern growth has been the increase in the fraction of output paid to compensate inventors for the fruits of their labor.

1. INTRODUCTION

The past century has been marked by extremely rapid increases in standards of living. Measured GDP per capita is perhaps ten times higher in the United States today than 125 years earlier, and with a mismeasurement of growth of one percentage point per year, the factor could easily be more than thirty.

Also remarkable is the relatively brief span of history during which this rapid growth has occurred. Conservative estimates suggest that humans were already distinguishable from other primates 1 million years ago. Imagine placing a time line corresponding to this million year period along the length of a football field. On this time line, humans were hunters and gatherers until the agricultural revolution, perhaps 10,000 years ago — that is, for the first 99 yards of the field. The height of the Roman empire occurs only 7 inches from the rightmost goal line, and the Industrial Revolution begins less than one inch from the field's end. Large, sustained increases in standards of living have occurred during a relatively short time — equivalent to the width of a golf ball resting at the end of a football field.

This paper combines an idea-based theory of growth in which people are a key input into the production of new ideas with a model of endogenous fertility and mortality in order to analyze these remarkable facts. The model features the classic race between technological progress and the Malthusian diminishing returns associated with a limited supply of land. A key requirement for technological progress to be able to win this race is the presence of increasing returns to accumulable factors. With this assumption, the internal dynamics provided by the model are able to produce thousands of years of virtually no sustained growth in standards of living despite increases in both technology and population, followed by the emergence of rapid growth and a demographic transition. It is this event — the emergence of large, sustained increases in standards of living — that we will refer to as an industrial revolution. While this is a slightly broader use of the term than is standard, I will ask the reader's indulgence; for the time scale considered in this paper, this is the key revolution that needs to be explained.

The last part of the paper explores an additional factor that is often assigned an important role in the emergence of modern growth by economic historians including North and Thomas (1973), Rosenberg and Birdzell (1986), Jones (1988), and Mokyr (1990): changes in institutions such as property rights. While the modeling of institutions in this paper is extremely rudimentary and reduced-form, quantitative analysis of the model assigns a major role to such changes in explaining growth over the very long run. A counterfactual experiment at the end of the paper suggests, for example, that absent the large improvements in innovation-promoting institutions measured to have occurred in the 20th century, the Industrial Revolution would have been delayed by more than 300 years.

This project builds on a number of recent studies of growth over the very long run, including Lee (1988), Becker, Murphy and Tamura (1990), Kremer (1993), Goodfriend and McDermott (1995), Acemoglu and Zilibotti (1997), Tamura (1998), Lucas (1998), Galor and Weil (1998), and Hansen and Prescott (1998). Following Lee (1988) and Kremer (1993), the link between population and the discovery of new ideas plays a critical role.[1] As in the human capital-driven models of Becker et al. (1990), Tamura (1998), Lucas (1998), and Galor and Weil (1998), fertility behavior is governed by utility maximization. Common to most of these papers and to this one is a Malthusian building block: a fixed supply of land that generates decreasing returns to scale when technology is held constant.

This paper differs from the existing literature in several ways. It is the first paper to place the population and ideas channel at the center of a complete, unified model of both population and standards of living over the entire course of history, including the distant past, the industrial revolution, the demographic transition, and modern growth. Second, the paper is the most explicit and successful attempt to date to match the dramatic change in the time path of living standards that occurred in the last two hundred years. Finally, the paper emphasizes the role of changes in institutions that promote innovation, and especially their importance for understanding previous cycles of growth and stagnation as well as the ultimate triumph of modern economic growth that we know as the Industrial Revolution.

The remainder of the paper is organized as follows. Section 2 presents the basic model. Section 3 analyzes the model's dynamics and discusses how it generates a demographic transition. Section 4 presents a summary of the facts the model should address, explains how parameter values are obtained, and exhibits the basic simulation of the model. Section 5 conducts the quantitative analysis. Section 6 discusses some of the results and implications, and Section 7 concludes.

---

[1] This link has been emphasized by Simon (1986) and Romer (1990), among others.

2. THE MODEL

2.1. People

We begin by describing an environment in which fertility is chosen in a utility maximizing framework, in the tradition of Becker (1960), Razin and Ben-Zion (1975), and Becker and Barro (1988). The economy consists of N_t identical individuals, where $t = 0, 1, 2, ...$ indexes time. Each individual obtains utility from consumption c_t and from the number of children b_t produced by the individual in period t, according to

$$u(c_t, b_t) = (1-\mu)\frac{\tilde{c}_t^{1-\gamma}}{1-\gamma} + \mu\frac{\tilde{b}_t^{1-\eta}}{1-\eta}, \tag{1}$$

where $\tilde{c}_t \equiv c_t - \bar{c}$ and $\tilde{b}_t \equiv b_t - \bar{b}$. The parameter $\bar{c} > 0$ denotes the subsistence level of consumption in this economy, and the parameter $\bar{b} \geq 0$ is related to the long-run rate of fertility, as we will see shortly.

We assume $0 < \mu < 1$, $0 < \gamma < 1$, and $0 < \eta < 1$. These parameter restrictions ensure that the elasticity of substitution between consumption $\tilde{c}$ and children $\tilde{b}$ is always greater than one. This simple assumption will play an important role in generating the demographic transition.[2]

Individuals are each endowed with one unit of labor per period, which they can use to obtain consumption or to produce children. Let ℓ_t denote the amount of time the individual spends working, and let w_t denote the wage earned per unit of time worked. The technology for producing children is straightforward: each unit of time spent producing children leads to $\alpha > \bar{b}$ births.

The individual's optimization problem at each time t is given by

$$\max_{c_t, b_t, \ell_t} u(c_t - \bar{c}, b_t - \bar{b}) \tag{2}$$

subject to

$$c_t = w_t \ell_t \tag{3}$$

and

$$b_t = \alpha(1 - \ell_t), \tag{4}$$

[2]We do not necessarily require $\eta < 1$. Let $z \equiv \frac{1-\mu}{\mu}\frac{\tilde{c}^{1-\gamma}}{\tilde{b}^{1-\eta}}$. Then the elasticity of substitution between $\tilde{c}$ and $\tilde{b}$ is given by $\frac{1+z}{\gamma+\eta z}$. It is constant when $\gamma = \eta$ and takes the usual value $1/\gamma$.

taking w_t as given. The fact that this optimization problem is static simplifies the analysis. This fact can be derived from a more general dynamic optimization problem under two assumptions. First, we assume that utility depends on the flow of births rather than on the stock of children. Second, we assume that the probability of death faced by an individual depends on aggregate per capita consumption, which individuals take as given. With these assumptions, the more standard dynamic optimization problem reduces to the sequence of static problems given above.

2.2. Production of the Consumption Good

The consumption good in this economy is produced using labor L_Y, land T, and a stock of ideas A. Total output of this consumption good, denoted Y, is given by

$$Y_t = A_t^{\sigma} L_{Yt}^{\beta} T_t^{1-\beta} \epsilon_t, \tag{5}$$

where $\sigma > 0$ and $0 < \beta < 1$, and ϵ_t is an exogenous productivity shock. This production function is assumed to exhibit constant returns to scale to the rivalrous inputs labor and land, and therefore increasing returns to labor, land, and knowledge taken together. As in Romer (1990), this assumption reflects a key property possessed by knowledge. Knowledge is nonrivalrous and can therefore be used at any scale of production without having to be reinvented. The amount of land in this economy is fixed and normalized so that $T = 1$.

2.3. Dynamics: Production of Ideas and People

The dynamics of this economy arise from two sources. First, people today produce knowledge that makes it easier to produce consumption goods in the future. As above, A_t denotes the stock of ideas at the start of period t. Therefore, $\Delta A_{t+1} \equiv A_{t+1} - A_t$ is the number of new ideas discovered during period t. In this economy, people produce new ideas according to

$$\Delta A_{t+1} = \delta L_{At}^{\lambda} A_t^{\phi}, \tag{6}$$

where L_A is the number of people engaged in producing ideas and $\delta > 0$, $\lambda > 0$ and $\phi < 1$ are assumed. The production of ideas is modeled very much like the production of any other good. Just as a larger labor force produces more widgets, a larger number of researchers produce more ideas.

As in Jones (1995), the parameter λ allows for diminishing returns to increasing the number of researchers at a point in time, a way to capture duplication in idea creation. The parameter ϕ allows the productivity of research to be either an increasing ($\phi > 0$) or a decreasing ($\phi < 0$) function of the stock of ideas that have been previously discovered.

The second source of dynamics in the model is demography. Between two periods, the change in the population is equal to the number of births minus the number of deaths:

$$\Delta N_{t+1} = b_t N_t - d_t N_t \equiv n_t N_t, \quad N_0 > 0. \tag{7}$$

The number of births per capita b_t is determined by the fertility behavior of individuals, discussed above. The mortality rate d_t is assumed to be a function of the average level of per capita consumption relative to subsistence, a useful summary measure of the technological capability of the economy as well as a measure that likely reflects the sensitivity of the population to disease and natural disasters. The mortality rate is given by

$$d_t(c_t/\bar{c}) = f(c_t/\bar{c} - 1) + \bar{d}, \tag{8}$$

where $f(\cdot)$ is some decreasing function such that $f(0) \geq 1 + \alpha$ and $f(\infty) = 0$. As consumption rises, the mortality rate falls. As per capita consumption falls to the subsistence level, everyone in the population dies. This characteristic implicitly defines what we mean by subsistence. Notice also that $\bar{d} \geq 0$ denotes the mortality rate in an economy with infinitely large consumption.

2.4. The Allocation of Labor and Factor Payments

Three rivalrous inputs to production exist in this economy: labor used to produce goods, labor used to produce ideas, and land. As a simplification, we assume land is not owned by anyone; i.e. it is treated as an external factor. Next, we define an exogenous variable $\pi_t \in [0, 1]$ as the fraction of total production of the consumption good that is paid to compensate inventive effort in period t. We think of π_t as capturing one aspect of the institutional structure of the economy: in some periods, the institutions encourage the production of new ideas by devoting a large amount of resources to this endeavor, while in others, the production of new ideas may be discouraged by institutions that limit the extent to which inventors can be compensated. Historically, such institutions have included support for research from monarchs or patrons, prizes, and the awarding of temporary monopoly power through patents. The relationship between this variable and intellectual property rights is clear, but the relationship extends to other property rights as well. For example, the value of a patent obviously depends on the enforcement of property rights more generally.

With this definition, payments to labor in the idea sector are

$$w_{At}L_{At} = \pi_t Y_t, \tag{9}$$

and payments to labor in the consumption sector are

$$w_{Yt}L_{Yt} = (1 - \pi_t)Y_t, \tag{10}$$

where w_A and w_Y are the wages paid per unit of labor in the two sectors. These wages will be equated in equilibrium by the free flow of labor between the two sectors.

The resource constraint for labor in this economy is

$$L_{Yt} + L_{At} = L_t \equiv \ell_t N_t. \tag{11}$$

2.5. Equilibrium

The setup of the economy is now complete and we can define the equilibrium.

Definition: A *static equilibrium* in this economy in period t is a collection of allocations and prices $(c_t, \ell_t, Y_t, L_t, L_{At}, L_{Yt}, b_t, w_t, w_{At}, w_{Yt})$ such that, given values of the state variables A_t, N_t, π_t, and ϵ_t, (i) the choice variables c_t, b_t, and ℓ_t solve the representative individual's maximization problem, (ii) people are indifferent between spending their time producing goods and ideas, i.e. $w_t = w_{At} = w_{Yt}$, and (iii) the resource constraint (11) is satisfied.

Definition: A *dynamic equilibrium* in this economy is a sequence of static equilibrium allocations for $t = 0, 1, 2, ...$, together with sequences for $\{A_t, N_t, d_t, n_t\}_{t=0}^{\infty}$, such that, given an exogenous sequence of shocks $\{\pi_t, \epsilon_t\}_{t=0}^{\infty}$ and given the initial conditions A_0 and N_0, the evolution of the economy satisfies the laws of motion in equations (6), (7), and (8) and the constraints $A_t \geq 0$ and $N_t \geq 0$.

Solving for the equilibrium is straightforward. The individual's maximization problem yields the first order condition

$$\frac{u_{\tilde{b}}}{u_{\tilde{c}}} = \frac{w_t}{\alpha}. \tag{12}$$

This says that an individual must be indifferent at the optimum between spending a little more time working and spending a little more time producing children.

With the preferences given by equation (1), this first order condition implies

$$\tilde{b}_t = \left(\frac{\alpha\mu}{1-\mu} \frac{\tilde{c}_t^{\gamma}}{w_t} \right)^{1/\eta}. \tag{13}$$

Along a balanced growth path in this model, $\tilde{c}_t$ and w_t will grow at the same rate. The assumption that $0 < \gamma < 1$, then, is what allows this model to exhibit a fertility transition, i.e. a situation in which fertility eventually declines as the wage rate rises.

This can be seen formally by noting that equations (3) and (4) imply a second relationship $c_t = w_t(1 - b_t/\alpha)$. Substituting this expression for c_t into (13) one gets an implicit expression for $\tilde{b}_t$ as a function of the wage w_t. Differentiating this expression, the sign of $\frac{d\tilde{b}}{dw}$ is the same as the sign of

$$\frac{\bar{c}}{\gamma w_t} - \left(\frac{1}{\gamma} - 1\right)\left(1 - \frac{b_t}{\alpha}\right).$$

The traditional income and substitution effects are reflected in the second term. As the wage goes up, the income effect leads individuals to increase both consumption and fertility. The substitution effect, on the other hand, leads people to substitute toward consumption and away from fertility: the discovery of new ideas raises the productivity of labor at producing consumption, but the technology for producing children is unchanged. If $\gamma < 1$, then the substitution effect dominates, while if $\gamma > 1$, the income effect dominates. As usual, if $\gamma = 1$, i.e. with log utility, these two effects offset. A third effect not traditionally present is reflected in the first term: as the wage rises, the subsistence consumption level which the consumer is required to purchase gets cheaper, leading consumers to have more after-subsistence income to spend on both more children and more consumption. This effect disappears as the wage gets large. The assumption that $0 < \gamma < 1$, then, leads the subsistence effect to dominate for small values of the wage and the substitution effect to dominate for large values of the wage, producing one component of the demographic transition: fertility rises and then falls as the wage rate rises.

The allocation of labor between the two sectors is even more straightforward. Define $s_t = L_{At}/L_t$ as the fraction of the economy's labor that works to produce new ideas. Equating w_{At} and w_{Yt} in equations (9) and (10) leads immediately to

$$s_t = \pi_t. \tag{14}$$

That is, not only is π_t the fraction of the economy's output devoted to compensate inventors, in equilibrium it is also the fraction of the economy's labor devoted to searching for new ideas. The remaining fraction $1 - \pi_t$ is engaged in producing goods.

Based on these conditions, the following proposition (proved in Appendix A.1) establishes a simple condition under which an interior static equilibrium exists and is unique.

Proposition: Let $a_t \equiv (1 - \pi_t)^\beta A_t^\sigma \epsilon_t / N_t^{1-\beta}$ be a measure of productivity in this economy. Assume this measure of productivity is sufficiently large that $(\bar{c}/a_t)^{1/\beta} < 1 - \bar{b}/\alpha$. Then, there exists a unique interior static equilibrium $\ell^*(a_t)$, $w^*(a_t)$, $c^*(a_t)$, and $b^*(a_t)$.

The technical condition in the proposition is needed for an interior solution. In the case in which this condition is just violated, $(\bar{c}/a_t)^{1/\beta} = 1 - \bar{b}/\alpha = \ell^*$. The population is so large relative to the technology level that diminishing returns to land reduces the wage leading to $c^* = \bar{c}$ and $b^* = \bar{b}$. Given the mortality function in equation (8), everyone in the economy would die in that period, and the population would be zero from then on.

3. DYNAMICS AND STABILITY

To see how the static equilibrium evolves over time, we proceed as follows. First, we shut down the shocks in the model: we assume for the moment that $\pi_t = \pi \in (0, 1)$ and $\epsilon_t = 1$ for all t. Then, we characterize the equilibrium along a steady-state balanced growth path. Finally, we explore the dynamics and the stability properties of this path.

3.1. Balanced Growth

A balanced growth path is a situation in which all variables grow at constant geometric rates (possibly zero). We will look for a balanced growth path in which ℓ, s, b, and d are constant. To characterize the balanced growth path of this economy, begin with the production function for new ideas, equation (6). This equation implies

$$\frac{\Delta A_{t+1}}{A_t} = \delta \frac{L_{At}^{\lambda}}{A_t^{1-\phi}}. \tag{15}$$

Along a balanced growth path, the left-hand-side of this equation is constant by definition, so the right-hand-side must also be constant. Since ℓ and s are constant, this is true when

$$G_A = G_{L_A}^{\frac{\lambda}{1-\phi}} = G_N^{\frac{\lambda}{1-\phi}}, \tag{16}$$

where G_z is defined as the gross growth rate of any variable z along a balanced growth path, i.e. $G_z \equiv \frac{z_{t+1}}{z_t}$. The economy can exhibit balanced growth only if $\phi < 1$.

Since $c_t = w_t \ell_t$ and $w_t = a_t / \ell_t^{1-\beta}$, consumption is given by $c_t = a_t \ell_t^{\beta}$. Also, note that a_t is proportional to $A_t^{\sigma} / N_t^{1-\beta}$. The constancy of ℓ along a balanced growth path then implies

$$G_c = G_w = G_a = \frac{G_A^{\sigma}}{G_N^{1-\beta}} = G_N^{\theta}, \tag{17}$$

where $\theta \equiv \frac{\lambda\sigma}{1-\phi} - (1-\beta)$. As the ratio in equation (17) indicates, there is a race in the model between technical progress and the diminishing returns implied by a fixed supply of land (holding the stock of ideas constant). The assumption of $\theta > 0$, which we now make, ensures that this race can be won by technical progress and makes sustained exponential growth in per capita income possible.

The assumption of $\theta > 0$ implies that the model is characterized by increasing returns to accumulable factors. For example, suppose that the production of ideas is homogeneous of degree one, so that $\lambda + \phi = 1$. It is easy to show that $\theta > 0$ then requires $\sigma + \beta > 1$. Recall that the nonrivalry of ideas motivated the assumption of constant returns to land and labor and increasing returns to land, labor, and ideas together — i.e. $\sigma > 0$. We require the stronger assumption that there are increasing returns to ideas and labor, holding land constant. That is, the increasing returns implied by nonrivalry must be sufficiently strong.

As in several recent papers, the growth rate of per capita income and consumption along the balanced growth path is proportional to the rate of population growth. Notice that

$$\frac{\Delta N_{t+1}}{N_t} = n(a_t) = b(a_t) - d(a_t), \tag{18}$$

where we abuse notation somewhat in writing $d(a_t)$ for $d(c(a_t)/\bar{c})$. In addition, $b(a_t) = \alpha(1 - \ell(a_t))$, so that a constant ℓ requires constant birth and mortality rates along the balanced growth path. Under what condition will these rates be constant? Recall that the first order condition for the individual's optimization problem in equation (13) leads her to set the excess birth rate $\tilde{b}_t$ proportional to $(\tilde{c}_t^{\gamma}/w_t)^{1/\eta}$. Therefore, we need $\tilde{c}_t^{\gamma}/w_t$ to be constant. Along a balanced growth path, however, c_t and w_t grow at the same rate. This implies that a balanced growth path occurs only asymptotically as c_t and w_t go to infinity, and the demographic transition effects associated with $\gamma < 1$ apply. As this happens, $\tilde{b}_t$ approaches zero so that b_t approaches $\bar{b}$. In addition, the mortality rate approaches $\bar{d}$. We assume that $\bar{b} \geq \bar{d}$ so that

$$G_N = 1 + \bar{b} - \bar{d} \geq 1. \tag{19}$$

Applying this result to equation (17), we see that along the balanced growth path,

$$G_c = G_w = G_a = G_{Y/N} = (1 + \bar{b} - \bar{d})^{\theta}. \tag{20}$$

Notice that increasing returns is not sufficient for positive per capita growth along the balanced growth path. If $\bar{b} = \bar{d}$, then there is no population growth in the long run and the balanced growth path has zero per capita growth.

One must also be careful with the asymptotic nature of this result. The balanced growth path in this model is an asymptotic result that applies only as a_t goes to infinity. For example, even if the balanced growth path has zero (geometric) per capita growth, the growth rate of per capita income will be positive in every period and the level of per capita income will go to infinity.

3.2. The Demographic Transition

A necessary prelude to characterizing the stability of the balanced growth path and the nature of the transition dynamics of the model is an analysis of population growth and the demographic transition. Recall that $n(a) = b(a) - d(a)$. We examine $b(a)$ and $d(a)$ in turn.

First, $b(a) = \alpha(1 - \ell(a))$. From our discussion of b earlier, it should not be surprising that b is a humped-shaped function of the wage, and hence of productivity a. This pattern, and the implied $\ell(a)$ schedule, are shown in

FIGURE 1. Working Time, Fertility, and Mortality

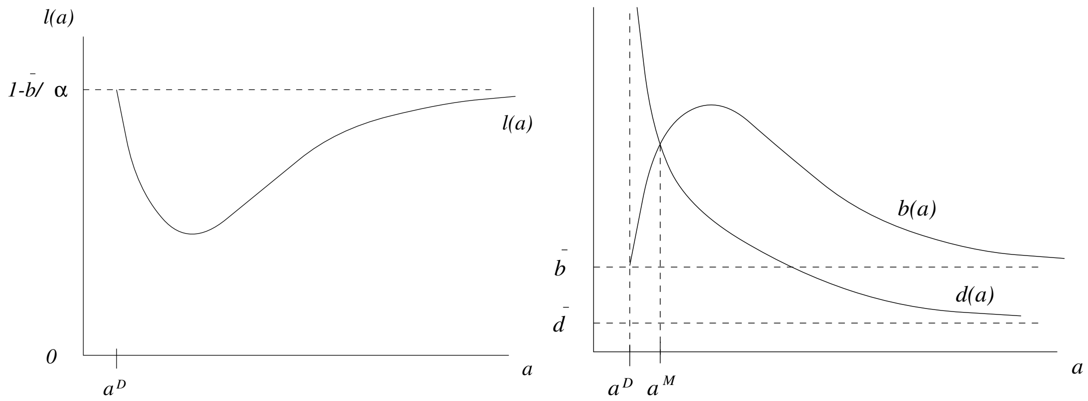

Figure 1.[3] Also shown in this figure is a $d(a)$ schedule; it is easy to show that $c(a)$ is a monotonic function, so that $d(a)$ has the same general shape as $d(c)$.

Finally, provided the function $f(\cdot)$ is restricted appropriately, we can now characterize $n(a) \equiv b(a) - d(a)$ as shown in Figure 2. The population growth rate is zero when productivity is a^M. It increases as a function of a to a level greater than $\bar{b} - \bar{d}$ (at least for a range of parameters values), and then declines to its balanced growth path level as a goes to infinity. The demographic transition is apparent in both Figures 1 and 2.

[3]To be more precise, $\ell(a)$ is the solution of the following nonlinear equation, obtained by combining (A.1) and $w = a/\ell^{1-\beta}$:

$$F(\ell) \equiv a(\alpha(1-\ell) - \bar{b})^{\eta} - \frac{\alpha\mu}{1-\mu}(a\ell^{\beta} - \bar{c})^{\gamma}\ell^{1-\beta} = 0.$$

Totally differentiating both sides of this equation with respect to a and ℓ, one sees that

$$sign\left\{\frac{d\ell}{da}\right\} = sign\left\{1 - \frac{\gamma}{1-\frac{\bar{c}}{c}}\right\}, \tag{21}$$

where $c = a\ell(a)^{\beta}$. The conditions that $\tilde{c} > 0$ and $\tilde{b} > 0$ limit the range of values that $\ell(a)$ can take to $(\bar{c}/a)^{1/\beta} < \ell < 1 - \bar{b}/\alpha$. When $a = a^D \equiv \frac{\bar{c}}{(1-\bar{b}/\alpha)^{\beta}}$, this range shrinks to the single point at which $\ell = (\bar{c}/a^D)^{1/\beta} = 1 - \bar{b}/\alpha$. On the other hand, we have already shown that $\lim_{a\to\infty} \ell(a) = 1 - \bar{b}/\alpha$. These endpoint conditions, together with the conditions on the slope given by equation (21) imply that the solution $\ell(a)$ has the shape given in Figure 1.

FIGURE 2. Population Growth as a Function of Productivity

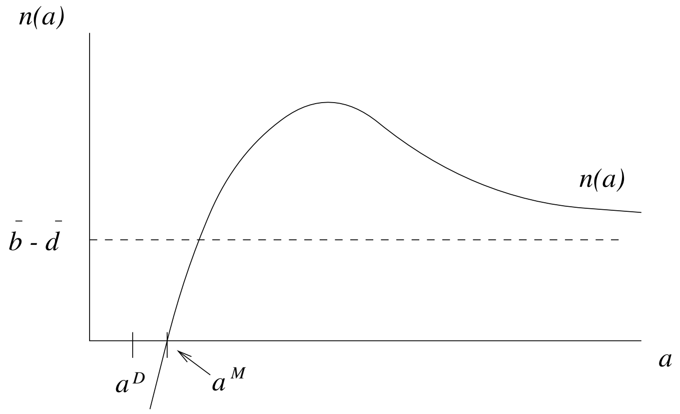

This general picture describes the classic version of the demographic transition. As summarized by Cohen (1995) and Easterlin (1996), the demographic transition consists of two phases. In the first, called a mortality revolution, mortality rates fall sharply, while birth rates either remain relatively constant or perhaps even rise slightly. The result is an increase in the population growth rate. The second phase is the fertility revolution, characterized by a birth rate that now falls more quickly than the relatively low but still declining mortality rate. The result is a decline in the population growth rate.[4]

3.3. Stability and Transition Dynamics

At a technical level, the balanced growth path in this model is globally stable; this result is documented in Appendix A.2. The transition dynamics themselves turn out to be quite rich. In particular, under the crucial assumption of increasing returns to accumulable factors ($\theta > 0$), the general

[4]The existing papers studying very long-run growth that have included a model of the demographic transition — Becker et al. (1990), Tamura (1998), Lucas (1998), and Galor and Weil (1998) — do so purely through a fertility transition that occurs as individuals begin to trade off quantity for quality. The models are set up in an overlapping generations context so that mortality is unaffected by technological progress. Galor and Weil (1996) generate a demographic transition through a difference in the endowments of men and women and a shift in comparative advantage. See Galor and Weil (1999) for an overview of several different theories of the demographic transition.

pattern is for growth rates of both population and standards of living to first increase and then to decrease, following the pattern displayed in Figure 2.

The intuition for this result is as follows. Initially, the economy begins with a small number of people who are not particularly good at producing ideas. Because new ideas are discovered infrequently, their main effect is to produce a temporary rise in consumption and a long-term rise in the size of the population — the standard Malthusian result. However, the increase in the size of the population means that new ideas are discovered with increasing frequency, which gives rise to the acceleration in growth rates emphasized by Lee (1988) and Kremer (1993). Under the increasing returns assumption, per capita consumption rises, gradually leading to the demographic transition, which slows population growth rates and therefore the growth rate of per capita consumption as well.

In contrast to many endogenous growth models which emphasize constant returns to accumulable factors, this model emphasizes the importance of increasing returns to factors that can be accumulated (including labor). The nonrivalry of ideas is an important feature in generating increasing returns, but it is not sufficient, due to the presence of land as a fixed factor. For the model to generate an industrial revolution endogenously, the nonrivalry of ideas must be sufficiently strong so as to overcome the diminishing returns implied by the fixed supply of land.

4. QUANTITATIVE ANALYSIS

4.1. The "Facts"

The model developed in the previous section will be analyzed quantitatively to help us understand growth over the very long run in both population and per capita income. First, however, we pause to present the "facts" about these two variables.

Cohen (1995) assembles data on world population from a number of studies conducted during the last forty years and provides a brief overview of the data. McEvedy and Jones (1978), the main source used by Kremer (1993), appears to be the most thorough study, and I will rely on Kremer's data, as reported in Table 1.

It is useful to appreciate both the extremely low rate of population growth over most of history as well as the time scale over which this rate operates. For example, using Kremer's collection of world population data, the rate of population growth, measured as the average annual change in log popula-

TABLE 1.
Population Data

Year	Population (Millions)	Average Annual Growth Rate over Preceding Period
-25000	3.34	...
-10000	4	0.000012
-5000	5	0.000045
-4000	7	0.000336
-3000	14	0.000693
-2000	27	0.000657
-1000	50	0.000616
-500	100	0.001386
-200	150	0.001352
0	170	0.000626
200	190	0.000556
400	190	0.000000
600	200	0.000256
800	220	0.000477
1000	265	0.000931
1100	320	0.001886
1200	360	0.001178
1300	360	0.000000
1400	350	-0.000282
1500	425	0.001942
1600	545	0.002487
1700	610	0.001127
1800	900	0.003889
1900	1625	0.005909
2000	5333	0.011884

Note: The levels of population are taken from Kremer (1993), who in turn takes his data from various sources. The population growth rate is computed as the average annual change in the natural log of population over the preceding interval. Two changes relative to Kremer are made. First, the year 1 A.D. is set equal to the year 0. Second, the population in 1990 is used for the population in the year 2000. These changes are made so that the period length in the model can be set equal to 25 years. The growth rates for a few periods are slightly different from those in Kremer because he reports growth rates from his underlying sources rather than based on the levels themselves.

tion, was only 0.0000072 between 1 million B.C. and 1 A.D. Nevertheless, over this period, the level of population increased by a factor of 1360: from 0.125 million people in 1 million B.C. to 170 million people in 1 A.D. A second key fact about population growth apparent in the table, emphasized by Kremer (1993), is that the rate of population growth is itself generally increasing over time. This is true not only in recent centuries but also dating back to our earliest data.

Data on per capita GDP or per capita consumption are much harder to come by, but the collection of evidence seems to support the following stylized picture: there was relatively little net increase in standards of living over most of history, say prior to the year 1500. Since then, per capita growth has risen, and levels of per capita income are now substantially higher than they were prior to 1500.

For example, Maddison (1982) estimates zero per capita income growth in Europe between 500 and 1500. Lee (1980) finds that the real wage in England in 1800 was nearly unchanged from its level in 1300; Hansen and Prescott (1998) make use of some new data assembled by Gregory Clark to reach a similar conclusion. Jevons (1896) uses detailed wage records from Athens in 328 B.C. to argue that wages in ancient Greece were roughly the same as those in Britain in the 15th century.[5] Schoenhof (1903) draws on price schedules covering more than 1000 items compiled by Diocletian in the year 302 A.D. to conclude that wages in ancient Rome were at least as high as those in France in 1790.

Levels of world GDP per capita can be constructed from Tables 1 and 2 in Lucas (1998). Such calculations imply levels (in 1985 international dollars) of $619 in 1750, $731 in 1850, $1764 in 1950, and $4257 in 1990. Alternatively, Maddison (1995) reports an estimate of $565 (in 1990 dollars) for the year 1500 and $5145 in 1992. DeLong (1998) reports values ranging from $115 to $512 in 1500 and $5204 in 1990, depending on whether or not an admittedly-coarse correction is made for quality change. Pritchett (1997) makes a case that per capita income was generally close to a lower bound of $250 for most of history.

Growth rates of world per capita GDP can be also be computed from these sources, and the results from Lucas and Maddison are in rough agreement. According to the numbers from Lucas (1998), the average annual growth

[5]Curiously, Clark (1940, p. 164ff) takes this calculation further to argue that the same statement is true of "modern" Britain, i.e. apparently in the 1920s or 1930s.

rate was 0.17 percent from 1750 to 1850, 0.88 percent from 1850 to 1950, and 2.20 percent from 1950 to 1990. Using decadal averages, annual world per capita GDP growth peaked in the 1960s at 3.10 percent per year before falling to 2.12 percent in the 1970s and 1.31 percent in the 1980s.

4.2. Parameter Choices

To simulate the model, values for 19 parameters are required. We will fix some of the parameter values ahead of time and then estimate others to fit the population data as well as possible. One issue that arises immediately is the distinction between consumption and GDP. This distinction is not present in the model, and we will choose for convenience to match up c in the model with data on per capita GDP.

In general, parameter values are picked to be both somewhat reasonable and to give the model its best chance of matching qualitatively the features of the data. We do not claim, of course, that this is the true model that generated the underlying data in the world. Rather, the goal is to produce *one* model that can generate the kind of patterns observed in the data, and to see what we learn from this model. For this reason, parameter choices will be discussed only briefly, as summarized in Table 2. The parameter $\bar{c}$ is set equal to 50, measured in 1990 dollars. If per capita consumption were to fall to $\bar{c} = 50$, everyone in the economy would die immediately. The parameter N_0 is set equal to 3.34, corresponding to the world population (in millions) in the year 25000 B.C., the first year of our simulation.

The parameter β is set equal to 2/3, so that the land share in an economy with perfect competition and property rights would be 1/3. This is the value chosen by Kremer (1993). The parameter α corresponds to the maximum birth rate at any instant in time, and we set the value of this parameter to 1/10. Easterlin (1996) and Livi-Bacci (1997, p. 7) report that maximum birth rates over history are about 0.05. With $\alpha = 1/10$, this birth rate occurs when one half of the individual's labor endowment is devoted entirely to raising children.

The parameters $\bar{b}$ and $\bar{d}$ correspond to the asymptotic birth rate and mortality rate in the model (that is, as consumption goes to infinity). How many kids would people like to have when consumption is infinite? We assume $\bar{b} = 0$, although one could make a case that this number should be positive. We also assume that the mortality rate goes to zero, so that people eventually can live forever. These assumptions imply that population growth goes to zero as consumption gets large.

TABLE 2.
Baseline Parameter Values

Parameter	Value	Comment
$\bar{c}$	50	Death with probability 1 at $c = \bar{c}$
N_0	3.34	Population (in millions) in year 25000 B.C.
β	2/3	Land share = 1/3 (Kremer)
α	1/10	Maximum fertility rate per period
$\bar{b}$	0	Long-run fertility rate
$\bar{d}$	0	Long-run mortality rate
σ	1	Unidentified
λ	3/4	Duplication of research
Period Length	25	Period length of 25 years
A_0	648.45	To produce $n_0 = 0$ in 25000 B.C.
ω_1	18.220	
ω_2	-0.082	Mortality function
ω_3	0.804	
μ	.9844	
γ	.8141	Utility function
η	.0250	

Note: Parameters of preferences and the birth and death technologies are estimated to fit birth and mortality rates, as described in the text.

For the elasticity of output with respect to new ideas, we set $\sigma = 1$. There always exists a value of ϕ consistent with any value of $\sigma > 0$ such that the model produces observationally equivalent results for population, consumption, and total factor productivity A^σ. These two parameters could be distinguished with data on the stock of ideas, but absent this data, they cannot be distinguished. We set $\sigma = 1$ so that we measure ideas in units of total factor productivity. The parameter ϕ, then, is conditional on this value, and would change (in a predictable fashion) for other values of σ.

The parameter λ is set equal to 3/4. If the population were instantaneously doubled, one suspects that the number of new ideas discovered would increase by less than a factor of two because the same idea would likely be discovered multiple times. This suggests an elasticity less than one. Choosing a specific value for λ is more difficult. Jones and Williams (2000) suggest that a value of 3/4 seems reasonable based on estimates of social rates of return. In the simulations below, this value produces plausible results.[6]

We assume a period length of 25 years, so that data on population are only observed infrequently.[7] The initial condition A_0 is chosen so that the population growth rate in the first 25-year period is equal to 0, given all of the values for the other parameters and assuming the initial value of π is vanishingly small. This leads to a value of $A_0 = 648.45$.

The remaining parameters which need values are related to the mortality function and the utility function. Recall that the mortality function is given by

$$d_t(z_t) = f(z_t) + \bar{d}, \quad z_t \equiv c_t/\bar{c} - 1. \tag{22}$$

We assume that $f(z)$ is the reciprocal of a polynomial: $f(z) = 1/(\omega_1 z^{\omega_2} + \omega_3 z)$. These polynomial parameters, together with the parameters of the utility function are estimated to fit the relationship between population growth rates and per capita consumption as well as possible, as described in Appendix A.3. Figure 3 plots the birth and mortality functions, together with the population growth rates that these parameter choices imply. In this and subsequent figures, "circles" will be used to denote data, corresponding in this figure to population growth and consumption. This simple fertility model performs well. The model generates a demographic transition and broadly matches the features of the population and consumption data.

4.3. Simulating the Model without Shocks

Figures 4 and 5 report the results from simulating the model with the parameter choices reported in Table 2 with a constant innovation share and no productivity shocks. For this simulation, we set $\pi = .005$, $\delta = .554$ and $\phi = 1/2$. These parameter values will be discussed further in the next section; for the moment, just take them as an example.

[6] Alternative estimates are available. Kremer (1993) estimates a value of λ greater than one, at 1.2. From the standpoint of this paper, a problem with this estimate is that Kremer's regression specification omits a potentially important variable: the innovation share. To that extent that the rise in population is correlated with an increase in the innovation share over time, this omitted variable bias will lead to an overestimate of λ. At the other end of the spectrum, Kortum (1992) estimates a value of λ of only 0.2 using data on productivity and R&D in U.S. industries. We discuss the robustness of the results to changes in λ below.

[7] As the fertility and mortality parameters and the population growth rate itself are calibrated to annual rates, the law of motion for population used in the simulation of the model is $N_{t+1} = (1 + n(a_t))^{25} N_t$.

FIGURE 3. Birth Rates, Mortality Rates, and Population Growth

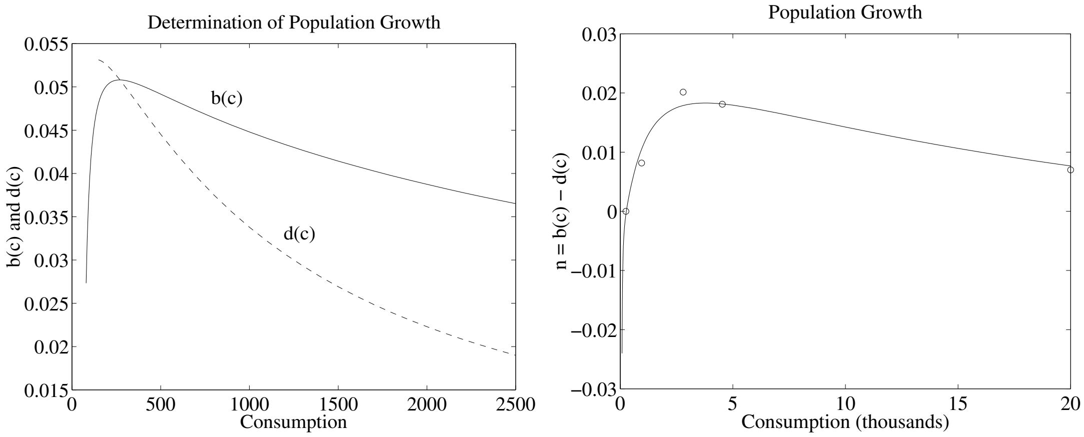

Note: The circles in the right graph represent observed data on population growth and per capita consumption. See Section A.3 and Table 8 in the appendix.

FIGURE 4. An Industrial Revolution

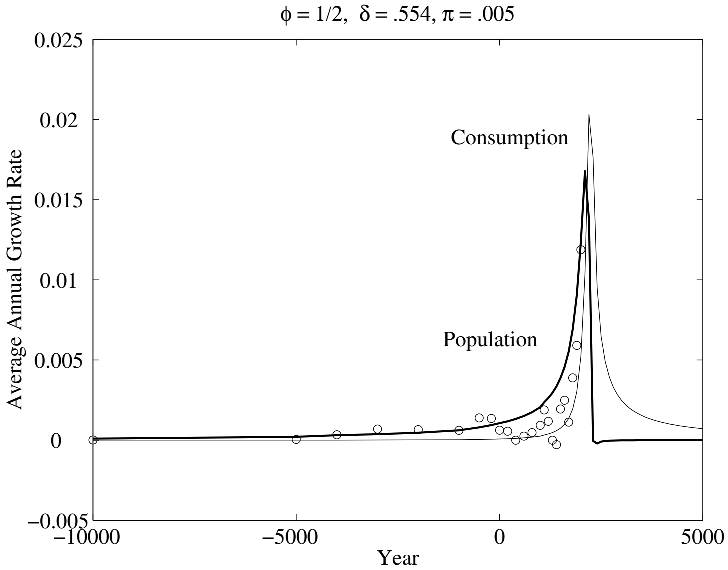

Note: The circles denote the average annual rates of population growth reported in Table 1.

FIGURE 5. Population and Per Capita Consumption

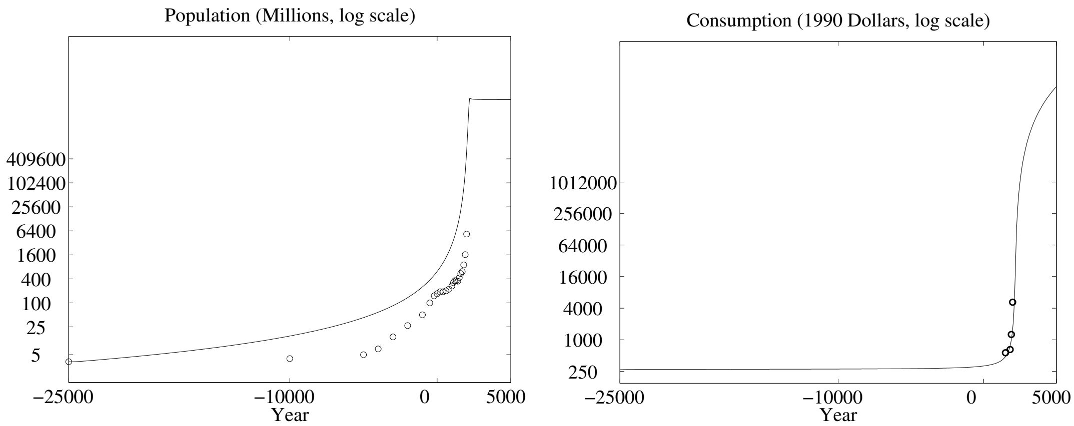

Note: The circles in the left figure denote the population growth rates in Table 1. The circles in the right panel are levels of per capita GDP from Maddison (1995) for 1500, 1820, 1900, and 1990.

Figure 4 reports average annual growth rates for the data on population (the circles) as well as the model's simulated growth rates for population and per capita consumption. Figure 5 displays the level of population and the level of per capita consumption, both in the simulation and for the data.

These figures illustrate that the internal dynamics of the model are able to replicate broadly the patterns observed in the data. The simulation exhibits thousands of years of very slow growth, followed by a sharp rise around the time of the Industrial Revolution. In levels, the model systematically overpredicts the level of population but does an excellent job of matching the data on per capita consumption. In particular, the level of consumption is stable for thousands of years before rising sharply with the Industrial Revolution.

The future predictions of the model are interesting even if they should not be taken seriously. With $\bar{b} = \bar{d} = 0$, the long run growth rate of the model is zero. Population growth falls to zero after the Industrial Revolution, as the rise in consumption generates a demographic transition. Consumption growth falls to zero gradually, but only after the level of per capita consumption is well on its way toward infinity. Implausibly, the level of population stabilizes at more than 400 billion!

5. CHANGES IN THE INNOVATION SHARE

With a constant innovation share, there exist values of π, ϕ, and δ that fit the broad patterns of the history of world population and per capita consumption. However, this simple model overpredicts the level of population systematically, and one may wonder where these parameter values come from. In this section, we add shocks to the allocation of resources to innovation π_t and to productivity so that the model fits the population data exactly. In the process, some useful findings will be uncovered.

The solution proceeds as follows. We solve for a sequence of innovation allocations π_t and productivity shocks ϵ_t, so that the model's simulated levels of population exactly match the actual levels. Both shocks are assumed to be constant during the entire interval between successive observations on the level of population (recall that a period is only 25 years and that we therefore observe the level of population infrequently). If a positive shock to the innovation share works, we shut off the productivity shock for that interval ($\epsilon = 1$). On the other hand, if the level of population declines or grows very slowly, it is possible that even a constant stock of ideas will overpredict the subsequent level of population. In this case, we set the innovation share π equal to zero to produce the constant stock of ideas over the interval. We then find the value of $\epsilon < 1$ such that the subsequent level of population is matched exactly. The solution method is provided in more detail in Appendix A.3.1.

Ideally, one would of course like to incorporate a richer pattern of shocks — for example, by allowing for productivity shocks even when the shock to the innovation share is positive. Given the limited nature of the data, however, it is not possible to identify the shocks in such a specification.[8]

Finally, the value of π in the 20th century is set equal to 0.05; some assumption like this is needed to pin down the value of δ. This value implies that the world economy in the 20th century spends roughly five percent of its output to compensate inventors and roughly five percent of its labor works to produce new ideas. Ratios of R&D to GDP in advanced countries are around three percent, but the definition of R&D implicit in this statistic is much narrower than the notion of ideas in the growth literature.

[8]It is possible to view the general procedure for estimating the shocks from a more formal econometric standpoint. Each π_t and ϵ_t can be thought of as a parameter to be estimated, together with μ, γ, and η. The discussion in the text explains how these parameters are estimated. Note that the number of parameters and data points is roughly the same in this exercise.

FIGURE 6. Innovation Shares for Different Values of ϕ

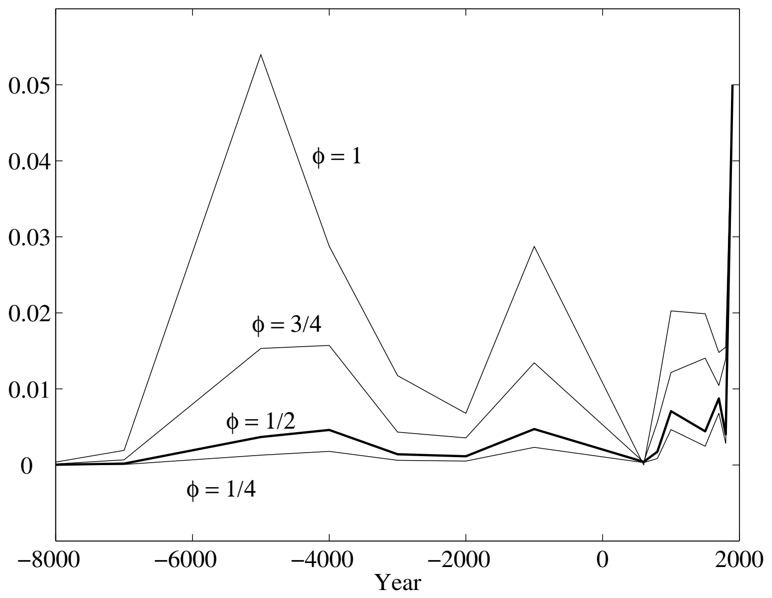

Note: The figure plots the sequence of innovation shares π_t needed to match the population data for four different values of ϕ. For each series, the value of π_{1900} is set equal to 0.05. To make the figure easier to read, periods when the innovation share is zero are ignored; see Table 3 for exact values. Also, the shocks for the years -25000 and -10000 are plotted at the years -8000 and -7000, respectively.

In any case, the nature of the results is not particularly sensitive to this parameter value.

Solving the model in this fashion, given the population data in Table 1, leads to the sequence of innovation shares graphed in Figure 6. The model is solved for four different values of ϕ, reflecting our uncertainty about this parameter value.

Regardless of the value of ϕ, each sequence yields the result that the value of the innovation share in the 20th century is significantly higher than in the preceding thousand years and is orders of magnitude higher than in 25,000 B.C. However, the different values of ϕ indicate very different innovation shares over the period in between. In particular, for large values of ϕ, the model suggests that around 5000 B.C., as much as five percent of output may have been devoted to compensate inventive effort, a level not reached again until the 20th century. Given what we know about the history of institutions and innovation, this seems implausibly high. Based on the results displayed in Figure 6, we will choose a value of $\phi = 1/2$ in the simulations that follow.[9] For most of the subsequent results — the main exception being the model's predictions for the future — the results are almost completely insensitive to the value of ϕ (in this range).

Table 3 reports the actual values of π_t and ϵ_t that are needed to fit the population data exactly. Each type of shock will be discussed in turn. Several remarks concerning the evolution of innovation shares are in order. First, with $\phi = 1/2$, the values of π_t suggest that innovation-promoting institutions have been getting better on average over the 25,000 year period. Just how much better can be seen in Table 4, which constructs some averages across several intervals.

According to the table, the number of ideas produced per year (measured in units of multifactor productivity) increased more than 110,000 times between the beginning of the simulation in 25,000 B.C. and the 20th century. A factor of 108 of this increase is due to the fact that there is a larger population available upon which to draw: more people produce more ideas. Interestingly, even with $\phi = 1/2$, only a factor of about 4 of this increase is associated with the rise in knowledge spillovers. Previous discoveries raise the productivity of research in the future, but this effect is estimated to be fairly small prior to the 20th century. The remaining factor of 245 is attributed to improvements in innovation-promoting institutions. That is, an important determinant of the technological advances that have made possible modern standards of living is the fact that the fraction of output devoted to compensating inventive effort has risen substantially. This effect is particularly acute when the 20th century is compared to the preceding several hundred years. Prior to 1900, increases in population and the innovation share were roughly of equal importance in contributing to the production of ideas. It is in the 20th century that the shift in the allocation of resources toward innovation becomes dominant.[10]

---

[9]The corresponding value of δ that leads to a value of $\pi = .05$ in the 20th century is 3.9923.

[10]This result depends somewhat on the value of λ. If $\lambda = 1$, toward the upper end of plausible values, the increase in L accounts for a factor of 515 while the contribution of π is reduced to around 50. This result maintains $\phi = 1/2$, however, which then implies too large an innovation share in 5000 B.C. Reducing ϕ to 1/8 raises the contribution of π to 150, leaving the contribution of L at 515. On the other hand, taking the very small estimate

TABLE 3.
Shocks: π_t and ϵ_t

Year	N_t	π_t	ϵ_t
-25000	3.34	0.00003	1
-10000	4	0.00017	1
-5000	5	0.00367	1
-4000	7	0.00460	1
-3000	14	0.00140	1
-2000	27	0.00114	1
-1000	50	0.00472	1
-500	100	0	0.972
-200	150	0	0.957
0	170	0	0.986
200	190	0	0.929
400	190	0	0.971
600	200	0.00040	1
800	220	0.00174	1
1000	265	0.00707	1
1100	320	0	0.820
1200	360	0	0.712
1300	360	0	0.681
1400	350	0	0.946
1500	425	0.00442	1
1600	545	0	0.773
1700	610	0.00875	1
1800	900	0.00402	1
1900	1625	0.05000	1
2000	5333	...	...

Note: These shocks are computed for the case of $\phi = 1/2$ and $\delta = 3.99232$.

TABLE 4.
Producing Ideas: $\Delta A_{t+1} = \delta \pi_t^\lambda L_t^\lambda A_t^\phi$

Interval			π_t^λ	L_t^λ	A_t^ϕ	New Ideas Per Year
-25000	to	-10000	1.0	1.0	1.0	1
-10000	to	-5000	3.5	1.1	1.0	4
-5000	to	0	26.9	1.4	1.1	39
0	to	1000	6.9	19.1	2.0	267
1000	to	1500	16.9	26.7	2.3	1027
1500	to	1900	38.8	38.2	2.6	3834
1900	to	2000	244.7	107.8	4.2	110467

Note: All series are normalized to one in the first interval. π_t^λ is computed as the average value of π over the interval (using the data in Table 3), raised to the power λ. L_t^λ and A_t^ϕ are calculated using the population and stock of ideas at the start of the interval.

This aspect of the simulation seems to be supported by historical evidence. Institutional changes throughout the world, such as the development of intellectual property rights, have surely raised the profits available to entrepreneurs, luring an increasing fraction of the population to search for new ideas.

The general rise in the allocation of resources to innovation occurs against a backdrop of fluctuations. According to the results, between the years 5000 B.C. and 1 A.D., the world population was especially active in generating ideas. The average value of π_t reaches a local peak during this interval at just under 1/2 of one percent, a level that is not exceeded systematically until recent times. Historically, this period marked the emergence of civilization in the form of cities. Key technological developments included writing, the beginning of scientific observation, the widespread use of metals, and dramatic improvements in transportation capabilities through the construction of ships and wagons. Whether or not these discoveries can be related to an improvement in the ability of entrepreneurs to earn returns is an open question that could be explored. This pattern fits with a view that the world during the ascendency of the Mesopotamian, Egyptian, and

of λ of 0.2 from Kortum (1992) implies contributions of 3.5 for L and more than 7500 for π for $\phi = 1/2$. The robust result seems to be that increases in both π and L are important, as conveyed in the table.

Greek civilizations was more productive at generating ideas than during the middle ages prior to the Enlightenment.

The productivity shocks may also be analyzed in this fashion. These shocks are shown in the last column of Table 3 and play the following role. The basic model contains forces that, at least until the demographic transition occurs, imply an ever increasing rate of population growth. In the data, in contrast, there are a number of periods during which population growth falls or even becomes negative. To account for these periods, we reduce productivity, which in turn reduces fertility and raises mortality.[11]

5.1. Simulation Results

With these shocks, our simulated economy is able to reproduce exactly the time series for world population. Figure 7 displays the actual and simulated data for population growth and consumption growth for the case of $\phi = 1/2$. After the year 2000, it is assumed that the allocation of resources to innovation π_t remains constant at its 20th century value and that no productivity shocks hit the economy. Summary statistics for this experiment are reported in Table 5.

That the simulation fits the population data exactly is in one sense not surprising — the shocks were chosen exactly for this purpose. What is remarkable, however, is that this fit is achieved with changes in the innovation share and productivity that appear reasonable given the historical record. For comparison, imagine the shocks that would be required for a standard neoclassical growth model to fit these same facts.

Two additional features of Figure 7 are worth noting. First, the time path of consumption growth broadly matches that outlined in our "facts" section: per capita consumption growth is quite close to zero until recent years, at which point it spikes up to nearly three percent per year. Second, regarding the future of population growth and consumption growth, both

[11]The most obvious place where such shocks are required is between the years 1100 and 1400, where the model requires the economy to run at only 3/4 of its full productivity potential. Significant shocks during this period include the genocidal Mongol invasions under Genghis Khan and his successors (in which perhaps a third of the population of China died) and the Black Death in Europe of the mid-14th century, which killed between a quarter and a third of the European population. The model also requires a surprisingly large productivity shock during the 17th century, in which productivity runs at only 79 percent. Notable shocks in this period include the Thirty Years' War in Europe, the Manchu conquest in China, and the continuation of the mass annihilation of Native Americans, particularly in Central and South America.

FIGURE 7. Population (Actual and Simulated) and Consumption

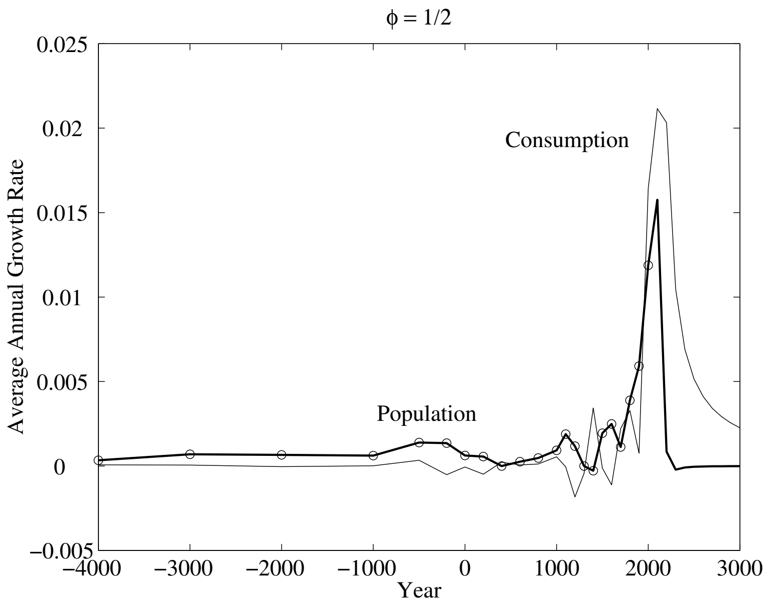

Note: The circles in the figure denote the data on population growth rates from Table 1. The lines denote simulated values.

peak sometime shortly after the year 2000 and then decline, eventually to zero since we have assumed that $\bar{b} = \bar{d} = 0$.

The simulation results in Table 5 show that the level of per capita consumption rises slightly from 25000 B.C. until the year 0. In contrast, the level of population rises from 3.34 million to 170 million, a 50-fold increase. This is the long-run Malthusian consequence of the improvements in technology shown in the first column.

The levels of population and consumption are shown in Figure 8. The apparent constancy of consumption for most of history suggested in the figure is an artifact of the time scale. Figure 9 plots the level of per capita consumption from 1000 B.C. until 1800 A.D. to illustrate this point. From an average level of \$270 throughout most of time, per capita consumption rises to \$300 in 1000 B.C. and reaches a local peak of about \$352 in 500 B.C. before falling back to \$298 by year 0. Reasonably large swings in consumption similar to this one continue through the year 1800, reflecting the impact of shocks to the innovation share and to productivity.

FIGURE 8. Population and Per Capita Consumption

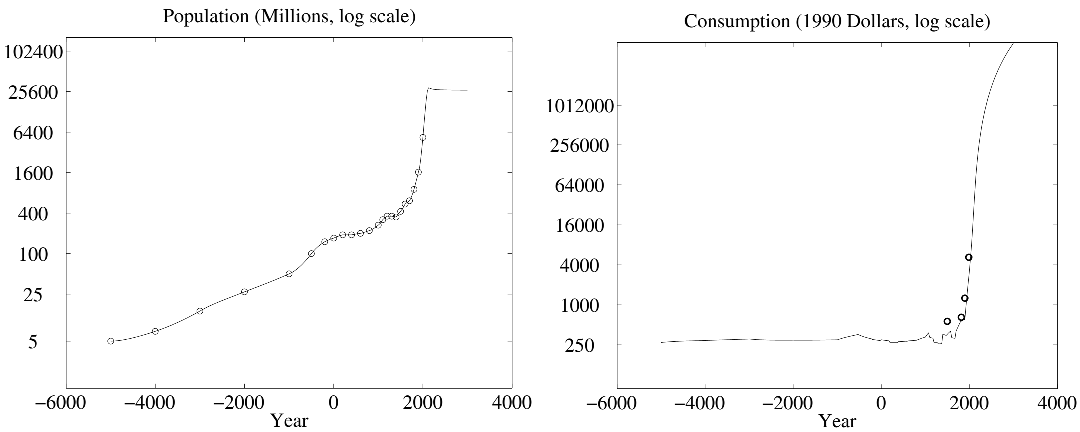

Note: The circles denote data; the lines represent simulated values.

FIGURE 9. Per Capita Consumption: 1000 B.C. to 1800 A.D.

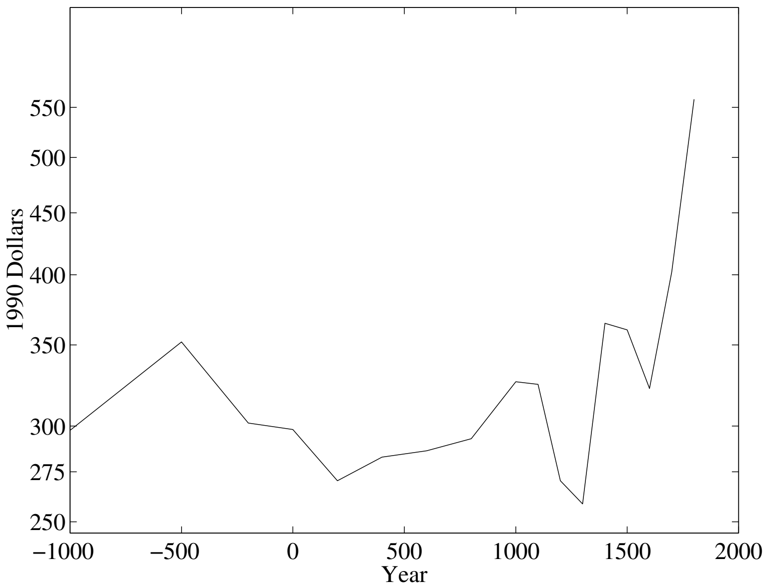

TABLE 5.
Simulation Results

Year	A	N	Average Growth	c	Average Growth
-25000	648	3.34	...	270	...
-10000	690	4	0.00001	271	0.00000
-5000	747	5	0.00004	272	0.00000
-500	2686	100	0.00067	352	0.00006
0	2686	170	0.00106	298	-0.00033
1000	3371	265	0.00044	326	0.00009
1500	4335	425	0.00094	360	0.00020
1600	5450	545	0.00249	322	-0.00111
1700	5450	610	0.00113	402	0.00221
1800	8429	900	0.00389	559	0.00329
1900	11366	1625	0.00591	603	0.00076
2000	74446	5333	0.01188	3116	0.01643
2100	831051	25778	0.01576	25855	0.02116
2200	6203610	28065	0.00085	197275	0.02032
2300	7554072	27471	-0.00021	562216	0.01047

Note: Simulation results assuming $\phi = 1/2$.

These fluctuations raise an interesting possibility. It has long been noted that several civilizations such as ancient Rome or China in the centuries following the previous millennium have witnessed spurts of growth and technological progress, only to succumb eventually to the end of per capita growth and even a decline in standards of living. How is this possible, and why has this fate been avoided in more recent times?

The model provides one possible explanation, consistent with the discussions of North and Thomas (1973), Jones (1988), Rosenberg and Birdzell (1986), Baumol (1990), and Mokyr (1990). The establishment of institutions that encourage the discovery and widespread use of new ideas can lead societies to outstrip Malthusian forces. However, the removal of these same institutions can allow the Malthusian forces to once again become dominant. The technological frontier must be constantly pushed forward in order to avoid the specter of diminishing returns associated with fixed resources. The history of "growth recurring," to use the evocative phrase of Eric Jones, may reflect the establishment and then elimination of institutions such as property rights in various civilizations.

Have we broken from this cycle? It is impossible to know, of course. However, the model suggests one insight related to this question. With small populations and few ideas, an increase in the fraction of resources devoted to innovation has a small effect on rates of discovery and therefore on standards of living: the Industrial Revolution is too far away for small changes to make much of a difference. However, as the Industrial Revolution gets closer in time, property rights themselves become more valuable: a shift in the allocation of resources toward innovation can lead the era of modern growth to occur sooner. The cumulative effect of thousands of years of discoveries has been to raise the world population and stock of ideas to levels at which the establishment of institutions such as property rights could lead to large and rapid improvements in technology and standards of living.

Returning to the broader pattern of population and consumption displayed in Figure 8, one sees a rapid rise in consumption around the year 2000, leading to the onset of the fertility transition. World population stabilizes at slightly more than 25 billion around the year 2100.[12]

These patterns of population and consumption growth can be seen more clearly in Figure 10, which focuses in on a 600 year period beginning with the 20th century. Consumption growth peaks at more than 2.5 percent around the year 2125. Recall that world per capita GDP growth seemed to peak in the 1960s at around 3 percent per year. While the timing of the peak is off (by more than a century), the magnitude is about right for $\phi = 1/2$. Population growth peaks in the year 2025, coming much closer to the actual peak in world population growth that seems to have occurred during the 1960s.

6. WAS THE INDUSTRIAL REVOLUTION INEVITABLE?

The different papers briefly reviewed in the Introduction that have looked into this question reach, sometimes implicitly and sometimes explicitly, different conclusions. On the one hand, in the models in most existing papers, especially Galor and Weil (1998) and Hansen and Prescott (1998), the dynamics in place from the beginning of time suggest that something like an industrial revolution was inevitable. On the other hand, the model in

[12]These values are quite sensitive to the value of ϕ. For example, with $\phi = 0$, consumption rises more gradually, delaying the onset of the decline in fertility. As a result, world population grows (implausibly) to more than 250 billion before stabilizing.

FIGURE 10. Growth: The 20th Century and Beyond

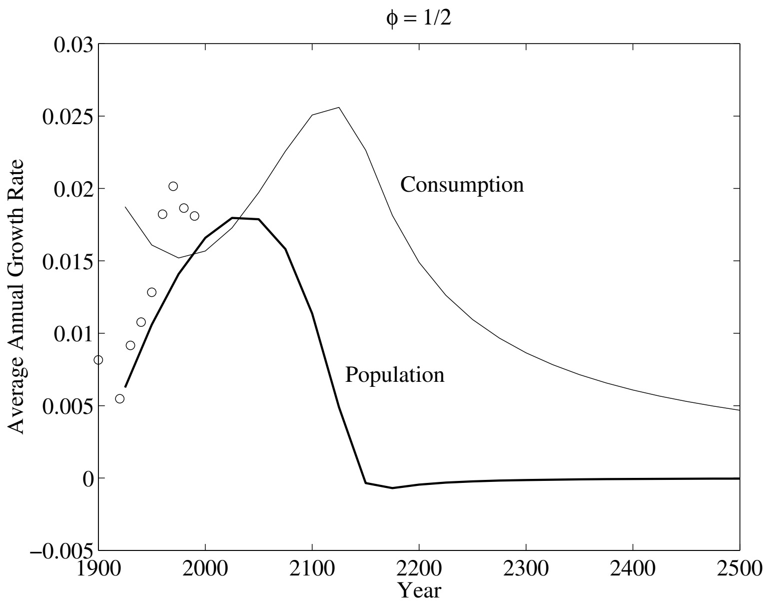

Note: The circles in the figure denote the data on population growth rates from Kremer (1993). The lines denote simulated values.

Lucas (1998) explicitly requires an exogenous shock to the rate of return to human capital accumulation in order to get the industrial revolution going. From a theoretical standpoint, one might imagine that this is an undesirable outcome, but from a historical standpoint — which might emphasize the development of property rights and the advent of science-based research — such a finding may be entirely appropriate.

The present paper is somewhere in between. A sensible working definition of an industrial revolution for this model is a substantial and rapid rise in both the level and growth rate of per capita consumption accompanied by a rise in population growth and followed by a demographic transition. Based on this definition, the model suggests that an industrial revolution was indeed inevitable, at least for the parameter values under consideration. This was apparent in Figure 4.

But was *the* Industrial Revolution inevitable? If by the Industrial Revolution we mean the onset of rapid population and per capita growth culmi-

FIGURE 11. Growth and Consumption: No 20th Century Shock

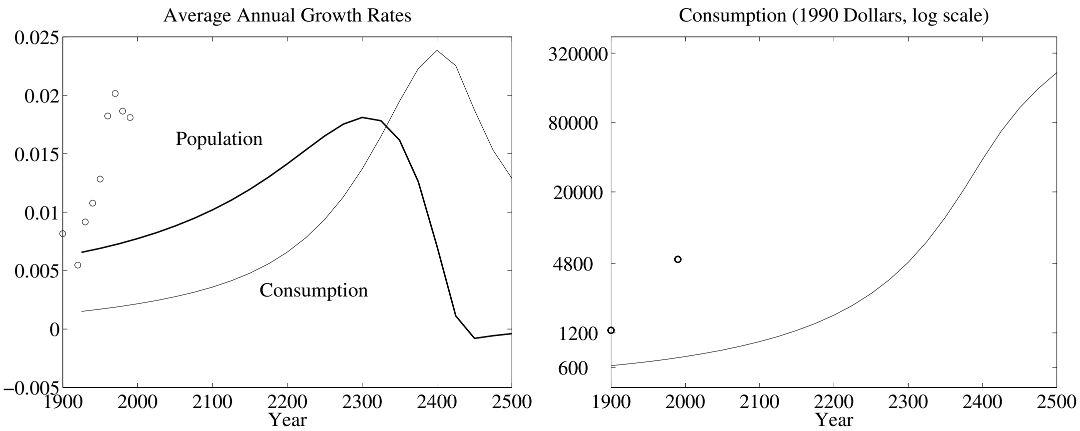

Note: The circles denote data from Kremer (1993) and Maddison (1995). The lines denote simulated values.

nating in the large increases in standards of living during the 20th century, then the answer turns out to be no.

To see this, consider the following counterfactual experiment. Suppose the large improvement in the 20th century in the institutions that promote innovation had never happened. Specifically, suppose π_t remained at its 19th century value forever. The simulation results for this case are reported in Figure 11. What we see from this experiment is that an industrial revolution does indeed occur, but it is delayed by more than 300 years. The shift in the allocation of resources toward innovation in the 20th century therefore played a critical role in the timing of the onset of modern growth, at least in the simulation.

7. CONCLUSION

First, a long-overdue caveat. The history of the world is obviously orders of magnitude more complicated than this simple model suggests. Many features of the history of growth are omitted and are undeniably important; a far from exhaustive list would include physical capital, human capital, endogenous property rights, a world of multiple regions, the diffusion of ideas, the shift from agriculture to industry, etc. The goal of this paper has been to provide one fully worked-out example of an economy that could have produced the basic patterns we think we observe. The hope is that some of the insights gained from this exercise shed light on the incredibly rich history of economic growth and will carry over into the next round of analysis.

The model of growth in this paper suggests the following basic story. A long time ago, the world population was relatively small and the productivity of this population at producing ideas was extremely low, in part because of the absence of institutions such as property rights. Once an idea was discovered, however, consumption and fertility rose, producing a rise in population growth. More people were then available to find new ideas, and the next new idea was discovered more quickly. In the model, this feedback leads to accelerating rates of population growth and consumption growth provided the aggregate production technology is characterized by increasing returns to accumulable factors.

Even in the absence of changes in the innovation share, this general feedback seems capable of producing something like an industrial revolution. However, the quantitative analysis suggests that changes in institutions to promote innovation are potentially extremely important. The rise and decline of institutions such as property rights could be responsible for the rise and decline of great civilizations in the past. And the establishment of innovation-promoting institutions in the 20th century appears to have played a critical role in generating the observed Industrial Revolution. Changes in such institutions have the potential to shift the onset of modern growth forward or backward in time. When the Industrial Revolution was hundreds or thousands of years away, such changes did not have large payoffs, providing one possible explanation for why the adoption of these institutions was so often temporary. Nevertheless, the cumulative effect of the discoveries in the past was to raise the population and stock of ideas to levels at which the adoption of institutions favorable to innovation could generate large increases in standards of living.

COLOPHON

Correspondence to author at Department of Economics, University of California, Berkeley, CA 94720-3880. Email: chad@econ.berkeley.edu. Web: http://elsa.berkeley.edu/~chad. I would like to thank Daron Acemoglu, Ken Arrow, William Baumol, Oded Galor, Avner Greif, Bob Hall, John Kennan, Michael Kremer, Lant Pritchett, Antonio Rangel, Paul Romer, Tom Sargent, Jon Temple, Alwyn Young, and seminar participants at Boston University, Columbia University, the Federal Reserve Bank of San Francisco, the Hebrew University Growth Conference, the Minneapolis Fed, the NBER Summer Institute, the Northwestern Summer Macro Workshop, NYU, Penn State, Stanford, and USC for their comments and suggestions. Sergei Morozov and Soojin Kim provided excellent research assistance. Financial support from the National Science Foundation (SBR-9818911) and the Sloan Foundation is gratefully acknowledged.

APPENDIX

A.1. EXISTENCE AND UNIQUENESS OF THE STATIC EQUILIBRIUM

Proof:

The first order condition in equation (13) can be combined with the two constraints in (3) and (4) to yield an implicit labor supply function $\ell(w)$:

$$\left(\alpha(1-\ell_t)-\bar{b}\right)^\eta = \frac{\alpha\mu}{1-\mu}\frac{(w_t\ell_t-\bar{c})^\gamma}{w_t}. \tag{A.1}$$

The wage is determined by the production function for consumption goods. Rewriting equation (10) with $L_{Yt} = (1-\pi_t)\ell_t N_t$ and recalling that $a_t \equiv (1-\pi_t)^\beta A_t^\sigma \epsilon_t / N_t^{1-\beta}$ yields

$$w_t = \frac{a_t}{\ell_t^{1-\beta}}. \tag{A.2}$$

Equations (A.1) and (A.2) can be combined to get a single nonlinear equation that characterizes the equilibrium value of ℓ_t. Dropping the time subscripts, define

$$F(\ell) \equiv (\alpha(1-\ell)-\bar{b})^\eta - \frac{\alpha\mu}{1-\mu}\frac{1}{a}(a\ell^\beta-\bar{c})^\gamma\ell^{1-\beta}.$$

Then the equilibrium satisfies $F(\ell^*) = 0$.

To see that there is a unique solution to this equation, first note that the Inada-type conditions on the utility function guarantee that a solution, if it exists, must satisfy $\tilde{c} > 0$ and $\tilde{b} > 0$. In terms of ℓ, these conditions imply that $\ell > (\bar{c}/a)^{1/\beta}$ and $\ell < 1-\bar{b}/\alpha$. Therefore, we require $(\bar{c}/a)^{1/\beta} < 1-\bar{b}/\alpha$ in order for a solution to exist. Given the definition of a, this puts restrictions on initial conditions.

Next, notice that $F(\bar{c}/a) > 0$ and $F(1-\bar{b}/\alpha) < 0$. Therefore, provided $F(\ell)$ is monotonically decreasing within this range, the solution is unique. The condition that $F'(\ell) < 0$ for $(\bar{c}/a)^{1/\beta} < \ell < 1-\bar{b}/\alpha$ is readily verified. Once $\ell^*(a)$ is determined, the remaining quantities in the proposition are given in a straightforward fashion from equations (A.2), (3), and (4). Q.E.D.

A.2. TRANSITION DYNAMICS

To analyze the stability properties of this economy, it is helpful to consider a simpler economy. Assume for this section that the production of new ideas depends on πN_t rather than on $L_{At} = \pi L_t$. This simplification does not change the qualitative nature of the transition dynamics.

Consider two state-like variables, the productivity variable a_t and a second variable $x_t \equiv \delta \pi^\lambda N_t^\lambda / A_t^{1-\phi}$. That is, x_t is the growth rate of A_t. The dynamics of x_t are given by

$$\frac{x_{t+1}}{x_t} = (1 + n(a_t))^\lambda / (1 + x_t)^{1-\phi}. \tag{A.3}$$

Since $a_t = (1-\pi)^\beta A_t^\sigma / N_t^{1-\beta}$ (again with $\epsilon_t = 1$), the dynamics of this state-like variable are given by

$$\frac{a_{t+1}}{a_t} = (1 + x_t)^\sigma / (1 + n(a_t))^{1-\beta}. \tag{A.4}$$

In addition, x_0 and a_0 are given by the initial conditions on A_0 and N_0. For interpretation, it is helpful to recall that c_t is monotonically related to a_t.

These two equations, together with the static equilibrium conditions that determine $n(a_t)$, completely characterize the dynamics of the economy. These dynamics can be examined in the discrete time version of a phase diagram. Notice that the balanced growth path occurs when $\Delta x_t = 0$ and $a_t = \infty$, so that the analysis of this system is slightly different from the traditional phase diagram analysis. Under the increasing returns assumption that $\theta > 0$, the $\Delta x_t = 0$ schedule lies "above" the $\Delta a_t = 0$ schedule, and the dynamics are characterized as in Figure 12.

As drawn in the figure, the dynamics are quite rich. The balanced growth path is globally stable: if the economy begins at any point such that $a_t > a^D$, it converges to the balanced growth path. However, we see that the growth rate of A_t will not generally be monotonic along the transition to the balanced growth path. It is natural to think of the economy starting from a point with a low a — low consumption — and a low x — slow technological progress. Then, the general pattern is for growth rates to rise and then fall as the economy approaches the balanced growth path.

It is also easy to see the importance of the increasing returns assumption that $\theta > 0$. If $\theta = 0$, so that the economy is characterized by constant returns to accumulable factors, then the $\Delta x_t = 0$ and $\Delta a_t = 0$ curves lie on top of each other. In this case, the dynamics of the economy move it

FIGURE 12. Phase Diagram

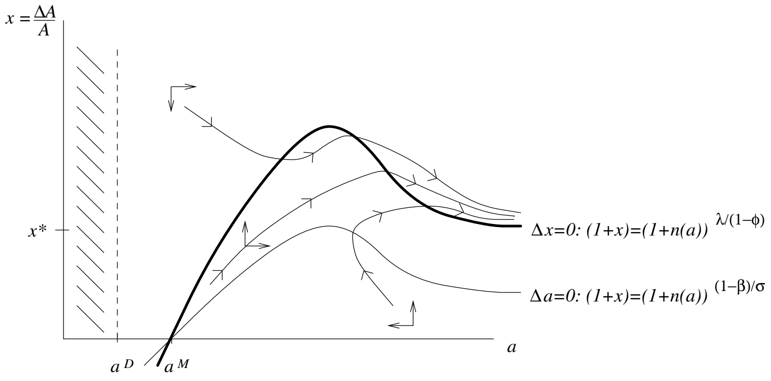

toward this curve, and the economy tends toward a situation in which a_t and x_t — and therefore per capita consumption — are constant.[1] If $\theta < 0$, the $\Delta a_t = 0$ schedule lies "above" the $\Delta x_t = 0$ schedule. In this case, there exists a globally stable steady state at the point $a_t = a^M$. Population growth (eventually) falls toward zero and consumption falls to some Malthusian-style subsistence level. In both of these cases, technological progress occurs forever, and the population grows to infinity. However, the key point is that technological progress does not translate into growth in per capita consumption, and the economy never experiences an industrial revolution.

Returning again to the $\theta > 0$ case, we are ready to consider what happens if shocks are added back into the system. Both the level of x and the level of a can jump as a result of shocks. A productivity shock that reduces ϵ causes a to decline. An increase in π causes the level of x to jump upward and the level of a to fall. However, apart from these jumps, the dynamics of the economy are still determined as in Figure 12.

Finally, consider the effect of an exogenous increase in mortality, such as the Black Death in 14th century Europe. Such a shock reduces N, causing a to jump to the right and x to jump down. The result is a rise in the level of the wage and the level of consumption in the short run, and a reduction in the rate of technological progress. As discussed by Lee (1988), population and per capita consumption can be negatively related in the short run even though they are positively related in the long run.

[1]One has to be careful here. Because time is discrete, the economy could cycle around such a point.

TABLE 6.
Observations on Mortality Rates

Per capita Consumption	"Data" Mortality Rate	Fitted Mortality Rate	Comments
100	.053	.053	
250	.05	.051	Livi-Bacci (1997)
800	.04	.038	
2000	.02	.022	Cohen (1995) for 1950-55
5000	.01	.011	Cohen (1995) for 1985-90
20000	.007	.003	Canada, 1989
100000	.001	.001	

Notes: The second and third data points are taken from rough guesses by Livi-Bacci (1997) (p. 7) that average mortality rates range as high as four to five percent and makes an educated guess that the mortality rate averaged something like four percent between 1 A.D. and 1750. The consumption numbers corresponding to these observations are simple guesses that seem plausible given the analysis of Pritchett (1997) on minimum income levels. The next two data points are taken from Cohen (1995, p. 68). Based on Maddison (1995), I assume that these years correspond to per capita GDPs of 2000 and 5000 dollars. Finally, the last number corresponds to the mortality rate in Canada in 1989 according to the World Bank (1991), Table 27. The mortality rate for the United States in that year was slightly higher, at 0.9 percent, while in Japan, Hong Kong, and Australia it was lower.

TABLE 7.
Parameter Estimates of the Mortality Function

Parameter	Coefficient	Standard Error
ω_1	18.220	(0.951)
ω_2	-0.082	(0.067)
ω_3	0.804	(0.153)
R^2	.996	

A.3. PARAMETER ESTIMATES OF MORTALITY AND PREFERENCES

We estimate the ω_i parameters in equation (22) using nonlinear least squares to fit the observations given in Table 6. The last column of numbers in Table 6 reports the fitted mortality rates; the coefficients themselves are reported in Table 7. The equation fits quite well, with an R^2 of 0.996.

TABLE 8.
Observations on Consumption and Population Growth

Per capita Consumption	"Data" Population Growth Rate	Fitted Population Growth Rate	Comment
250	0	-.0005	
948	.008164	.0105	Year 1875–1900
2792	.020151	.0179	Year 1960–1970
4533	.018101	.0181	Year 1980–1990
20000	.007	.0077	
50000	0	-.0004	

Given the mortality function $d(c)$, we turn to estimating the parameters related to fertility. As with death rates, we have very little information upon which to base our estimates of these parameters. The observations we draw on are given in Table 8 and describe population growth and consumption. The first and the last two observations in this table are rough guesses. Motivated in part by Pritchett (1997), we assume that the Malthusian consumption level at which births and deaths are equalized is something like two hundred fifty dollars (in 1990 dollars). The next to last observation corresponds roughly to per capita income and population growth in the richest countries today. The intermediate observations are taken from two sources. The consumption levels correspond to the world per capita GDP levels reported by Maddison (1995) in Table G-3.[2] The population growth rates, corresponding to the years reported in the comment, are taken from Kremer (1993).

Given the mortality function $d(c)$ that we have already estimated, and given $\alpha = .10$, we estimate μ, γ,and η to fit the population growth rate data in Table 8 as well as possible. Specifically, we estimate these parameters using nonlinear least squares to minimize the sum of squared deviations between the observed population growth rate and the model's predicted population growth rate at the given levels of consumption. The results of this estimation are reported in Table 9.

[2]Maddison does not report a value for 1875. I use the interpolated value from DeLong (1998).

TABLE 9.
Parameter Estimates of the Utility Function

Parameter	Coefficient	Standard Error
μ	.9844	(.0012)
γ	.8141	(.0164)
η	.0250	(.0299)
R^2	.986	

A.3.1. Solving for π_t and ϵ_t

Given the parameter values in Tables 2, 7, and 9, and given the population data in Table 1, we solve for the sequence of innovation shares $\{\pi_t\}$ and productivity shocks $\{\epsilon_t\}$ where $t = 0$ corresponds to the year 25000 B.C., and each unit increment to t corresponds to an increment of 25 years. The solution is obtained as follows:

1. We begin with an initial population, an initial stock of ideas, and an observation for population some periods later. Let $NumPeriods$ denote the number of periods between the two observations on population. For example, if the first period corresponds to the observation in the year -2000 and the next is the year -1000, we have 1000/25+1 = 41 periods.

2. Solve for the constant value of the shock π such that the dynamics of the model would lead population to grow from its level at the first observation to its level at the second observation after $NumPeriod$ periods, with a percentage error less than or equal to 10^{-8}. If such a value is found and is "small" in the sense that it does not involve passing through the entire demographic transition in one period, then we're done with this step. Set the productivity shock for period $NumPeriod$ equal to one since it is not needed.

3. With respect to the previous step, there are two things to note. First, there are occasionally multiple values of the π shock that will work. We choose the smallest value (so that we are on the pre-demographic transition side of the population growth schedule as much as possible). Second, for declines in the level of population, or for relatively small increases, it is possible that no "small" shock will work. In this case, set the property rights shock for the periods corresponding to $1 : (NumPeriods - 1)$

equal to zero, and solve for the reduction in productivity — the constant value of $\epsilon < 1$ — such that the simulation matches the level of population after $NumPeriod$ periods, with a percentage error less than or equal to 10^{-8}.

4. Advance to the next population observation and repeat this process, starting with step 1 above, until all population observations have been fit by the model.

REFERENCES

Acemoglu, Daron and Fabrizio Zilibotti, "Was Prometheus Unbound by Chance? Risk, Diversification, and Growth," *Journal of Political Economy*, August 1997, *105* (4), 709–751.

Baumol, William J., "Entrepreneurship: Productive, Unproductive, and Destructive," *Journal of Political Economy*, 1990, *98* (5), 893–921.

Becker, Gary S., "An Economic Analysis of Fertility," in "Demographic and Economic Change in Developed Countries," Princeton, NJ: Princeton University Press, 1960.

___ and Robert J. Barro, "A Reformulation of the Economic Theory of Fertility," *Quarterly Journal of Economics*, February 1988, *108* (1), 1–25.

___ , Kevin M. Murphy, and Robert Tamura, "Human Capital, Fertility and Economic Growth," *Journal of Political Economy*, October 1990, *98*, S12–S37.

Clark, Colin, *The Conditions of Economic Progress*, London: Macmillan and Co., 1940.

Cohen, Joel E., *How Many People Can the Earth Support?*, New York, NY: W.W. Norton, 1995.

DeLong, J. Bradford, "Estimating World GDP, One Million B.C. - Present," December 1998. U.C. Berkeley mimeo.

Easterlin, Richard A., *Growth Triumphant*, Ann Arbor: The University of Michigan Press, 1996.

Galor, Oded and David N. Weil, "The Gender Gap, Fertility, and Growth," *American Economic Review*, June 1996, *86* (3), 374–387.

___ and David Weil, "Population, Technology, and Growth: From the Malthusian Regime to the Demographic Transition," 1998. NBER Working Paper No. 6811.

—— and ——, "From Malthusian Stagnation to Modern Growth," *American Economic Association Papers and Proceedings*, May 1999, *89* (2), 150–154.

Goodfriend, Marvin and John McDermott, "Early Development," *American Economic Review*, March 1995, *85* (1), 116–133.

Hansen, Gary D. and Edward C. Prescott, "Malthus to Solow," 1998. NBER Working Paper No. 6858.

Jevons, F. B., "Some Ancient Greek Pay-Bills," *Economic Journal*, September 1896, *6* (23), 470–475.

Jones, Charles I., "R&D-Based Models of Economic Growth," *Journal of Political Economy*, August 1995, *103*, 759–784.

—— and John C. Williams, "Too Much of a Good Thing? The Economics of Investment in R&D," *Journal of Economic Growth*, March 2000, *5* (1), 65–85.

Jones, E. L., *Growth Recurring: Economic Change in World History*, New York: Oxford University Press, 1988.

Kortum, Samuel, "Inventions, R&D and Industry Growth," June 1992. Boston University mimeo.

Kremer, Michael, "Population Growth and Technological Change: One Million B.C. to 1990," *Quarterly Journal of Economics*, August 1993, *108* (4), 681–716.

Lee, Ronald D., "A Historical Perspective on Economic Aspects of the Population Explosion: The Case of Preindustrial England," in Richard A. Easterlin, ed., *Population and Economic Change in Developing Countries*, Chicago: University of Chicago Press, 1980.

——, "Induced Population Growth and Induced Technological Progress: Their Interaction in the Accelerating Stage," *Mathematical Population Studies*, 1988, *1* (3), 265–288.

Livi-Bacci, Massimo, *A Concise History of World Population*, Malden, MA: Blackwell Publishers, 1997.

Lucas, Robert E., "The Industrial Revolution: Past and Future," 1998. University of Chicago mimeo.

Maddison, Angus, *Phases of Capitalist Development*, New York: Oxford University Press, 1982.

——, *Monitoring the World Economy 1820-1992*, Paris: Organization for Economic Cooperation and Development, 1995.

McEvedy, Colin and Richard Jones, *Atlas of World Population History*, New York: Penguin, 1978.

Mokyr, Joel, *The Lever of Riches*, New York: Oxford University Press, 1990.

North, Douglass C. and Robert P. Thomas, *The Rise of the Western World*, Cambridge: Cambridge University Press, 1973.

Pritchett, Lant, "Divergence: Big Time," *Journal of Economic Perspectives*, Summer 1997, *11* (3), 3–17.

Razin, Assaf and Uri Ben-Zion, "An Intergenerational Model of Population Growth," *American Economic Review*, December 1975, *65* (5), 923–933.

Romer, Paul M., "Endogenous Technological Change," *Journal of Political Economy*, October 1990, *98* (5), S71–S102.

Rosenberg, Nathan and L. E. Birdzell, *How the West Grew Rich: The Economic Transformation of the Industrial World*, Basic Books, 1986.

Schoenhof, Jacob, "History of the Working Classes and of Industry in France," *Journal of Political Economy*, June 1903, *11* (3), 416–447.

Simon, Julian L., *Theory of Population and Economic Growth*, New York: Basil Blackwell, 1986.

Tamura, Robert, "From Agriculture to Industry: Human Capital and Specialization," 1998. Clemson University mimeo.

World Bank, *World Development Report, 1991*, Oxford University Press, 1991.