

Population and Ideas: A Theory of Endogenous Growth

CHARLES I. JONES

1. INTRODUCTION

Can exponential growth be sustained forever? How do we understand the exponential increase in per capita income observed over the last 150 years?

The growth literature provides a large number of candidate theories to address these questions, and such theories are nearly always constructed so as to generate a steady state, also known as a balanced growth path. That is, the growth rate of per capita income settles down eventually to a constant. In part, this reflects modeling convenience. However, it is also a desirable feature of any model that is going to fit some of the facts of growth. For example, as noted by Barro and Sala-i-Martin (1995, p. 34):

[O]ne reason to stick with the simpler framework that possesses a steady state is that the long-term experiences of the United States and some other developed countries indicate that per capita growth rates can be positive and trendless over long periods of time. . . . This empirical phenomenon suggests that a useful theory would predict that per capita growth rates approach constants in the long run; that is, the model would possess a steady state.

Clearly there are many examples of countries that display growth rates that are rising or falling for decades at a time. However, there are also examples of countries, such as the United States over the last 125 years, that exhibit positive growth for long periods with no noticeable trend. It seems reasonable, then, that a successful theory of growth should at least admit the possibility of steady-state growth.

Models with this property, however, are very special and require strong assumptions. One of these assumptions is that technical change, at least in the long

I thank Daron Acemoglu, Robert Barro, Robert Lucas, Antonio Rangel, Paul Romer, T. N. Srinivasan, John Williams, Alwyn Young, and seminar participants at the Chicago GSB, Cornell, the European Science Foundation Conference on Growth, Harvard, the Minneapolis Fed, MIT, the NBER Summer Institute, Princeton, UCSD, and Stanford for insightful comments. I am grateful to the Hoover Institution and the National Science Foundation (SBR-9818911) for financial support.

run, should not be capital augmenting. Another is the presence of a differential equation that is exactly linear in a sense we will define shortly.

Now that growth theorists understand the kind of assumptions that have to be made to generate sustained exponential growth, it is possible to construct a large number of models with different “engines” of growth, ranging from physical capital accumulation to human capital accumulation to the discovery of new ideas to population growth to various combinations of these factors. Indeed, the problem now confronting growth economists is how to choose among the abundance of competing explanations. Empirical work provides some guidance, but a number of difficulties such as the accurate measurement of ideas or human capital or even growth itself lead this research to be less than conclusive.

This chapter proposes a complementary approach to judging growth models by “raising the hurdle” to which our models aspire. Specifically, the suggestion is that a successful theory of economic growth should provide an intuitive and compelling justification for the crucial assumptions that are a requirement of such a theory. That “crucial” assumptions should be justified is a time-honored strategy for making progress in the growth literature. This kind of reasoning is discussed explicitly in the introduction in Solow (1956) and is partly responsible for the discovery of the neoclassical growth model.¹ Another example relates to the requirement that technical change should not be capital augmenting in the long run. At first, this seems like a very ad hoc assumption. However, research several decades ago by Kennedy (1964) and Drandakis and Phelps (1966) and more recently by Acemoglu (2001) explains how this can be the natural outcome in a model in which researchers choose the direction of technical progress.

Just as these previous authors made progress by questioning the justification for ad hoc but crucial assumptions, I propose that additional progress toward understanding long-run growth can be made by seeking a justification for the kind of linearity that is needed in models that generate sustained growth over long periods of time. Existing growth models fall short of this ideal, providing essentially no justification for why a key differential equation should be linear. This was surely appropriate when we were searching for the first several candidate explanations of long-run growth, but perhaps it is now time to ask more of our models.

The final result of any model that exhibits long-run growth is an equation of the form $\dot{y}/y = g$, where y is per capita income and $g > 0$ is a constant. Not surprisingly, then, the key to obtaining such a result is for the model to include a differential equation that is “linear” in a particular sense, as in

$$\dot{X} = \underline{\quad} X. \quad (1)$$

1. “All theory depends on assumptions that are not quite true. That is what makes it theory. The art of successful theorizing is to make the inevitable simplifying assumptions in such a way that the final results are not very sensitive. A ‘crucial’ assumption is one on which the conclusions do depend sensitively, and it is important that crucial assumptions be reasonably realistic. When the results of a theory seem to flow specifically from a special crucial assumption, if the assumption is dubious, the results are suspect. I wish to argue that something like this is true of the Harrod-Domar model of economic growth” (Solow, 1956, p. 65).

Growth models differ according to the way in which they label the X variable and the story they tell in order to fill in the blank.²

Much of the work in both new and old growth theory can be read as the search for the appropriate characterization of equation (1). For example, the original models in Solow (1956) and Swan (1956) without exogenous technical change focused our attention on the differential equation for capital accumulation. However, with diminishing returns to capital, that equation was less than linear, and there was no long-run growth in per capita income. When Solow and Swan added exogenous technical change in the form of an equation that was assumed to be linear, $\dot{A} = gA$, long-run growth emerged.

The so-called "AK" growth models departed from Solow and Swan by eliminating the diminishing returns to capital accumulation. Linearity in the accumulation of physical capital or human capital (or some combination of these two) became the engine of growth.³ Idea-based growth models by Romer (1990), Grossman and Helpman (1991), Aghion and Howitt (1992), and others returned the linearity to the differential equation for technological progress and filled the blank in equation (1) with resources devoted to research by profit-maximizing entrepreneurs.

Note that *exact* linearity of the key differential equation is critical to generating sustained exponential growth in the long run. If the exponent on X in equation (1) is slightly larger than one, then growth rates will explode over time, with the level of X (and hence income) becoming infinite in a finite amount of time. On the other hand, if the exponent on X is slightly less than one, then growth rates will fall to zero asymptotically. In other words, the growth theorist is in the strange situation of requiring a knife-edge restriction.

One can argue that too much emphasis is placed on exact linearity in the previous paragraph. The U.S. evidence suggests that trendless growth is possible for at least 125 years, at a rate of about 1.8 percent per year. Matching this kind of evidence requires a differential equation that is close to linear. For example, if the differential equation takes the form $\dot{y}_t = ay_t^\phi$, acceptable values for ϕ fall approximately into the range 0.95–1.05.⁴ If one wants balanced growth forever at a positive rate, exact linearity is a requirement. If one only desires to match the empirical evidence

2. This way of summarizing growth models is taken from Romer (1995). It is important to recognize, as documented by Mulligan and Sala-i-Martin (1993), that this linearity can be hidden in models with multiple state variables. Linearity is also an asymptotic requirement rather than something that must hold at all points in time, as pointed out by Jones and Manuelli (1990).

3. The models of Romer (1987), Lucas (1988), and Rebelo (1991) fit this category.

4. Integrating the differential equation in the text for $\phi \neq 1$ and calculating the average growth rate leads to

$$\bar{g}_T \equiv \frac{1}{T} (\log y_T - \log y_0) = \frac{1}{T} \frac{1}{1-\phi} \log[1 + g_0(1-\phi)T],$$

where g_0 is the growth rate \dot{y}/y at time 0 (corresponding here to the year 1870). Setting $\bar{g}_T = 0.018$ and $T = 125$, one can solve this equation for the value of ϕ associated with any initial value of g_0 . If $\phi < 1$, then growth rates are declining; a value of $g_0 = 0.019$ implies a value of $\phi = 0.952$ and a 1995 growth rate of $g_T = 0.0171$. What we know of U.S. history suggests that growth rates were rising prior to the 125 year period, and setting $g_0 = 0.017$ leads to a value of $\phi = 1.051$ and a 1995

for the United States, this requirement can be relaxed slightly. In either case, a crucial assumption of such a model is that it deliver a differential equation that is approximately linear.

A natural requirement of a successful theory of economic growth, then, is that it provide a compelling and intuitive justification for linearity; such an assumption is crucial to the result so we might require a good explanation for why it holds. On this basis, existing models are clearly deficient. The linearity in existing models is assumed *ad hoc*, with no motivation other than that we must have linearity somewhere to generate endogenous growth.

That a theory of endogenous growth assumes this kind of knife-edge linearity has long been known. Stiglitz (1970) and Cannon (1998) note that this requirement made growth theorists uncomfortable with models of endogenous growth in the 1960s. Solow (1994) appeals to this same criticism in arguing against recent models of endogenous growth. What is sometimes not sufficiently well appreciated is that any model that is going to generate sustained exponential growth requires such an assumption. A productive response to the criticism, then, is to provide justifications for our crucial assumptions.

This chapter develops a new theory of endogenous growth in which linearity is motivated from first principles. The process illustrates the potential gains from forcing ourselves to jump over a higher hurdle—the model has predictions that are different from those of other endogenous growth models. The first key ingredient of this model is endogenous fertility. At an intuitive level, the reason why endogenous fertility helps is straightforward. Consider a standard Solow-Swan model. With the labor force as a factor that cannot be accumulated endogenously, one has to look for a way—typically arbitrary—to eliminate the diminishing returns to physical capital. In contrast, with an endogenously accumulated labor force, both capital and labor are accumulable factors, and a standard constant-returns-to-scale setup can easily generate an endogenously growing economy.

However, endogenous fertility in a model with constant returns to scale in all production functions will not generate endogenous growth in *per capita* variables. This leads to the second key ingredient of the model: increasing returns to scale. Endogenous fertility leads to endogenous growth in the scale of the economy. Increasing returns to scale in the production function for aggregate output translates the endogenous growth in scale into endogenous growth in per capita output.

Research on idea-based growth models provides a justification for increasing returns that is based on first principles. At least since Shell (1966), Phelps (1968), and Nordhaus (1969), economists have recognized that the nonrivalry of knowledge implies that aggregate production is characterized by increasing returns to scale. This argument has been clarified and elevated to a very prominent place in our thinking about economic growth by Romer (1990). Ideas are nonrivalrous; they can be used at any scale of production after being produced only once. For example, consider the production of any new product, say the digital videodisc player or the latest worldwide web browser. Producing the very first unit may

growth rate of $g_T = 0.0191$. The growth rates in this last experiment are roughly consistent with those estimated by Ben-David and Papell (1995).

require considerable resources: The product must be invented or designed. However, once the product is invented, it never needs to be invented again, and the standard replication argument implies that subsequent production occurs with constant returns to scale. Including the production of the “idea,” or the design of the product, production is characterized by increasing returns. This property, rather than the assumption that the differential equation governing technological progress is linear, is the key contribution we need from the idea-based growth literature.⁵

This chapter builds on a number of earlier insights. Several papers in the 1960s contain key results that are further developed here. Phelps (1966) and Nordhaus (1969) present models in which the nonrivalry of knowledge leads to increasing returns and derive the result that long-run growth in per capita income is driven by exogenous population growth.⁶ Still, neither of these papers seems to know how seriously to take this prediction, with Nordhaus calling it a “peculiar result” (p. 23). Two years later, however, Phelps (1968, pp. 511–12) stresses the implications of population for growth:

One can hardly imagine, I think, how poor we would be today were it not for the rapid population growth of the past to which we owe the enormous number of technological advances enjoyed today. . . . If I could re-do the history of the world, halving population size each year from the beginning of time on some random basis, I would not do it for fear of losing Mozart in the process.

More recently, Jones (1995) modified the Romer (1990) model to eliminate the apparently counterfactual prediction that the growth rate of the economy is proportional to the size of the population. In the modified model, the growth rate of the economy depends on the growth rate of the population, as in the earlier models.⁷ Because the population growth rate is assumed to be exogenously given, however, the long-run growth rate of the economy is invariant to policy changes.⁸

Here, the population growth rate is endogenized, and policy changes can affect the long-run growth rate of the economy through their effects on fertility. However, as the channel through which policy affects growth is fertility, the nature of the

5. Alternative methods for introducing increasing returns to scale in the model, such as a Marshallian externality associated with capital accumulation, will also lead to endogenous growth. I focus on the idea-based theory of increasing returns because it can be motivated from first principles.

6. The learning-by-doing models of Arrow (1962) and Sheshinski (1967) also obtain this result.

7. A large body of related research includes Simon (1981), Judd (1985), Grossman and Helpman (1989), Kremer (1993), Raut and Srinivasan (1994), Kortum (1997), and Segerstrom (1998).

8. Subsequent papers by Dinopoulos and Thompson (1998), Peretto (1998), Young (1998), and Howitt (1999) have found clever ways to eliminate the effects of scale on growth in idea-based models without eliminating long-run policy effects. These models maintain linearity in the equation for technical progress but assume that the number of sectors grows exactly with population, so that research effort per sector does not grow. See Jones (1999) for a discussion of these issues.

effects of policy on long-run growth is often counter to conventional wisdom. For example, subsidies to R&D and capital accumulation, even though they may be welfare improving, will reduce long-run growth in the model.

Section 2 of the chapter presents an extremely simple growth model that illustrates the role of population growth and increasing returns. The basic claim in this section is that if one takes the historical presence of population growth as given, then one can understand growth in per capita income without introducing any arbitrary linearity into the model.

The remainder of the chapter then examines the deeper issue of how we can understand growth more generally, both in per capita terms and in population, as an endogenous phenomenon. Section 3 explores in detail the claim that endogenous fertility can provide the linearity needed to understand per capita growth. Section 4 develops the decentralized dynamic general equilibrium model in the context of "basic science." That is, the model is based on the assumption that not only are the ideas underlying growth nonrivalrous, they are pure public goods. This assumption is employed almost entirely because it simplifies the analysis considerably. Still, it may also be of independent interest. For example, it is sometimes conjectured that basic science should be modeled as an exogenous process, like exogenous technical progress in a Solow model. The analysis here suggests that insight is gained by moving beyond this view. Even if the ideas of basic science fall from above like apples from trees, the fertility channel and increasing returns are crucial: The number of Isaac Newtons depends on the size of the population that is available to sit under trees.

Section 5 explores the welfare properties of the model. Section 6 contains a general discussion of the model's predictions and discusses its interpretation. One point worth emphasizing from the beginning is that the model is best thought of as describing the OECD or even the world as a whole. Care is required when testing the model with a cross section of countries because countries share ideas.

2. THE ISAAC NEWTON GROWTH MODEL

The first claim in this chapter is that the growth in per capita income that has occurred in recent centuries can be understood without appealing to any extra linearity of the kind assumed in recent growth models. To make this claim, we construct an extremely simple toy economy and show how it exhibits per capita income growth.

There are two key ingredients that drive per capita growth in the toy model, both of which are readily justified. The first is population growth. For the moment, we simply take constant exogenous population growth as a given. Letting L_t represent the population or labor force at time t ,

$$\dot{L}_t/L_t = n > 0, \quad L_0 > 0 \text{ given.} \quad (2)$$

The second key ingredient is increasing returns to scale. Let Y_t be the quantity of a single consumption/output good produced, and let A_t be the stock of ideas

that the economy has discovered in the past. The production function in our toy model is

$$Y_t = A_t^\sigma L_{Yt}, \quad (3)$$

where L_Y is the number of people working to produce the output good and $\sigma > 0$ imposes the assumption of increasing returns to scale. Holding the stock of ideas constant, there are constant returns to scale: Doubling the quantity of rivalrous inputs (here only L_Y) will double output. Because ideas are nonrivalrous, the existing stock A can be used at any scale of production, leading to increasing returns in A and L_Y together.

Finally, we need a production function for ideas. This part of the model can be set up in a number of different ways. To keep the model simple, however, assume the following production function:

$$\dot{A}_t = \delta L_{At}, \quad A_0 > 0 \text{ given}, \quad (4)$$

where L_A is the number of people working to produce new ideas (the number of Isaac Newtons) and $\delta > 0$ represents the number of new ideas that each researcher discovers per unit of time.

The resource constraint for this economy is

$$L_{Yt} + L_{At} = L_t. \quad (5)$$

As part of our simplifying assumptions, we assume that a constant fraction s of the labor force works as researchers so that $L_{At} = sL_t$ and $L_{Yt} = (1-s)L_t$, with $0 < s < 1$. This is the only allocative decision that needs to be made in this simple model.

From the production function in equation (3), consumption (or output) per worker is given by

$$y_t \equiv Y_t/L_t = A_t^\sigma (1-s),$$

and therefore the growth of consumption per worker, g_y , is equal to σg_A , where $g_x \equiv \dot{x}/x$ for any variable x .

From the production function for ideas in equation (4),

$$\dot{A}_t/A_t = \delta s (L_t/A_t). \quad (6)$$

It is then easy to show that there exists a stable balanced growth path for this model where $g_A = n$. For example, in order for \dot{A}/A to be constant in equation (6), the ratio L/A must be constant. Therefore, the long-run per capita growth rate in this economy is given by

$$g_y = \sigma n. \quad (7)$$

This result nicely illustrates the central roles of population growth and increasing returns. Per capita growth is proportional to the rate of population growth, where the factor of proportionality measures the degree of increasing returns in the economy.

According to this model, sustained, long-run per capita growth results from population growth and increasing returns. The inherent nonrivalry of ideas means that the economy is characterized by increasing returns to scale. Economic growth occurs because the economy is repeatedly discovering newer and better ways to transform labor into consumption. However, the creation of new ideas by itself is not sufficient to generate sustained growth. For example, suppose an economy invents 100 new ideas every year. As a fraction of the (ever-evolving) existing stock of ideas, these 100 new ideas become smaller and smaller. Sustained growth requires that the number of new ideas itself grow exponentially over time. This in turn requires that the number of inventors of new ideas grow over time, which requires population growth.

If population growth is taken as given, this model suggests that per capita income growth is not a puzzle at all. More people means more Isaac Newtons and therefore more ideas. More ideas, because of nonrivalry, mean more per capita income. Therefore, population growth, combined with the increasing returns to scale associated with ideas, delivers sustained long-run growth.⁹

3. LINEARITY AND GROWTH

The previous section shows how two basic ingredients, population growth and increasing returns, can help us make sense of the presence of per capita income growth. In the remainder of the chapter, we consider a deeper question. How do we understand past and possibly future growth, both in per capita terms and in population itself, as an endogenous phenomenon?

Motivated by the discussion in the introduction, the answer must involve a differential equation that is linear. In this section, I argue that the law of motion for population is intimately tied to a linear differential equation in a way that the law of motion for physical capital or human capital or ideas is not.

9. The model clearly indicates that growth occurs because the effective resources devoted to producing new ideas increase over time. Jones (2002) applies a more general version of this model that incorporates both physical and human capital to uncover empirically the sources of twentieth-century U.S. economic growth. A key fact in the application is that resources devoted to research have increased for three reasons. In addition to basic population growth, the share of the labor force devoted to research and the educational attainment of the researchers have increased as well. I document that roughly 80 percent of postwar U.S. growth is due to increases in human capital investment rates and research intensity and only 20 percent is due to the general increase in population. However, the intensity effects cannot lead to sustained exponential growth—neither educational attainment nor the share of the labor force devoted to research can increase forever. So unless there is an ad hoc Lucas-style linearity in human capital accumulation, population growth remains the only possible source of long-run growth, leading that paper to predict that growth rates may slow considerably in the future. These results confirm that substantial progress in understanding growth can be made using the basic framework given earlier.

To understand this claim, imagine a world consisting of N_t (representative) individuals at time t . Each individual in this economy chooses to have a certain number of children, denoted by \tilde{n}_t . At each point in time, some constant, exogenously given fraction d of the population dies, as in the Blanchard (1985) constant-probability-of-death model. The law of motion for the aggregate population in a continuous time environment is then given by

$$\begin{aligned}\dot{N}_t &= (\tilde{n}_t - d)N_t, \\ &\equiv n_t N_t.\end{aligned}\tag{8}$$

Therefore, by deciding on the number of children to have, individuals choose the proportional rate of increase in the population. The linearity of the law of motion for population is a biological fact of nature: People reproduce in proportion to their number.

This is not to say that such an equation automatically delivers sustained exponential population growth. Indeed, n may depend on the aggregate state of the economy. In the model, for example, it will depend on the wage rate and other endogenous variables. In fact, in a model with decreasing returns to scale (e.g., because of a constant technology level and a fixed supply of land), a subsistence requirement for consumption, and endogenous fertility, one easily arrives at a Malthusian result in which the size of the population is asymptotically constant—people endogenously choose $\tilde{n} = d$, delivering zero population growth.

Instead, the point of this exercise is simply that thinking about fertility delivers an equation that is linear in a way that thinking about physical or human capital accumulation or the production function for knowledge does not. To see this more clearly, consider a very rough comparison of this equation to the key linear differential equation in other growth theories:

1. AK model: $\dot{K} = sK^\phi$
2. Lucas model: $\dot{h} = uh^\phi$
3. Romer model: $\dot{A} = H_A A^\phi$
4. Fertility model: $\dot{N} = (\tilde{n} - d)N^\phi$

Each of the models maintains the assumption that $\phi = 1$ (which may be viewed as an analytically useful approximation for the crucial assumption of $\phi \approx 1$).

What does it mean for these equations to be linear? Hold constant the control or choice variable of individual agents and consider whether doubling the state variable will double, more than double, or less than double the change in the state variable. For example, in the AK model, hold constant the saving rate chosen by individuals. What happens to net investment when the stock of capital in the economy is doubled? In a neoclassical model with the usual diminishing returns, net investment is less than doubled, so the neoclassical model is less than linear. The AK model, however, eliminates these diminishing returns through an ad hoc assumption. In the Lucas-style model, hold constant the fraction of time u that individuals spend accumulating skills and double the stock of human capital. If

a seventh grader and a high school graduate both go to school for 8 hours a day, does the high school graduate learn twice as much?

In the Romer-style model, let A denote the stock of ideas or designs and H_A be the total level of resources the economy devotes to research. Holding H_A constant, suppose we double the existing stock of ideas. What happens to the output of new ideas? A benchmark case of constant returns would be $\phi = 0$: The number of new ideas created by 100 units of research effort is independent of the total stock of ideas discovered in the past. One might suppose that $\phi > 0$ —the productivity of research is higher because of the discovery of calculus or the semiconductor. Or one might suppose that $\phi < 0$ —the most obvious ideas are discovered first, and it becomes more and more difficult to discover new ideas because of “fishing out.” What one sees from this example is that the case of $\phi \approx 1$ is clearly ad hoc. There is no intrinsic justification for linearity in the production function for new ideas.

In contrast, consider finally the fertility model. Hold constant the choice variable of individuals—the number of children per person, \bar{n} . What happens to the total number of offspring if we double the population? Of course the total number of offspring doubles. Linearity in the fertility equation results from the standard replication argument.¹⁰

One can endogenize the fertility rate \bar{n} by following the endogenous fertility literature associated with Dasgupta (1969), Pitchford (1972), Razin and Ben Zion (1975), and Becker and Barro (1988), among others. Individuals care not only about their own consumption, but also about the number of their descendants and the consumption of their descendants. This literature, especially Barro and Becker (1989), shows that the population can grow endogenously at a constant exponential rate in a neoclassical-style growth model.

This result depends in part on the production technology for children. Suppose

$$\bar{n} = bl^\psi, \quad (9)$$

where $0 \leq l \leq 1$ is the time an individual with a fixed labor endowment of one unit spends producing offspring, and $0 < \psi < 1$. The parameter $b > d$ (for “births”) represents the maximum number of children that an individual can have in a given period (i.e., if $l = 1$). With these properties, as we will see below, it is easy to get the result of positive, steady-state population growth. For example, all we need is that b be sufficiently large.

Now consider the possibility that instead of being a parameter, b depends directly on the state variables of the economy. For example, new ideas in health care might allow children to be produced with less labor effort; technological progress might increase b , although one might suspect that fertility remains bounded from above. What is critical, however, is that asymptotically b does not decrease with

10. One might wonder about a dependence of d on N , which could destroy the linearity. It seems most natural to think of the mortality rate d as depending on per capita consumption and on the technological sophistication of the economy. Increases in consumption or medical technology, for example, may reduce the mortality rate, perhaps leading it to asymptote to some constant level (perhaps even to zero). Incorporating these features into the model would not destroy the linearity of the fertility equation.

N . For example, to get sustained population growth, we must rule out a case like $b = N^{-\theta}$ with $\theta > 0$. Clearly, this would eliminate the linearity of the model.

Is it reasonable to believe that θ is not too far from zero? I think so. To see why, note first that people are a primary input into nearly all production functions. People are needed to produce output, people are needed to produce ideas, and people are needed to produce new people. At least so far, the AK assumption that machines by themselves can produce new machines does not seem tenable. With an exogenously given population, one needs an exact but ad hoc degree of increasing returns to scale to get constant returns to “reproducible” inputs. Thus, for example, the simple Romer equation given above requires a returns to scale of approximately two, which is difficult to justify. However, once population is itself an endogenously reproducible input, then a standard constant-returns-to-scale production function already exhibits constant returns to reproducible inputs; the two coincide so that the standard replication argument provides the key justification for linearity.

This is the situation that applies here. The standard constant-returns-to-scale benchmark in the production of offspring corresponds to $\theta = 0$. It is possible, of course, for θ to be substantially larger or smaller than zero, but *this* would require a departure from constant returns through some kind of arbitrary and difficult-to-justify external effect: As the population gets larger, why should the maximum number of children that an individual can produce decline?

4. THE DECENTRALIZED MODEL

The remainder of this chapter should be read as an extended example. We embed an endogenous fertility setup into an idea-based growth model and examine the kind of results that can arise. I have chosen a particular theory and made particular assumptions to get to the basic results easily. I will indicate in the appropriate places how the results generalize.

4.1. Preferences

One of the key insights of Barro (1974) was to think about utility-maximizing individuals who care not only about their own consumption but also about their children’s consumption. This reasoning was extended by Razin and Ben-Zion (1975) and Becker and Barro (1988) to model endogenous fertility: Parents also care about the number of children that they have, and there may be costs to increasing the number of offspring.

Following Becker and Barro (1988), we assume that the time s utility of the head of a dynastic family is given by

$$U_{0,s} = \int_s^{\infty} e^{-\rho(t-s)} u(c_t, \tilde{N}_{0,t}) dt, \quad (10)$$

where c_t is the consumption of a representative member of the dynastic family at time t , and $\rho > 0$ is the rate of time preference. $\tilde{N}_{0,t} \equiv N_t/N_0$ represents the size

of the dynastic family living at time t . Individuals live through their descendants, so that the death of an individual is not a remarkable event in that person's life. When an individual dies, her assets are divided evenly among the other members of the dynastic family.

With respect to the kernel of the utility function, it turns out to be convenient to assume

$$u(c_t, \tilde{N}_t) = \log c_t + \epsilon \log \tilde{N}_t, \quad (11)$$

where $\epsilon > 0$. Both the marginal utility of consumption and the marginal utility of progeny are positive but diminishing. The elasticity of substitution between consumption and progeny is one, as in Barro and Becker (1989). Within the class of utility functions with a constant elasticity of substitution between consumption and progeny, this unit elasticity guarantees that the dynastic approach is time consistent—choices made by the dynastic head of generation zero will be implemented by subsequent generations. This assumption also turns out to be required for the existence of a balanced growth path, as we will see shortly.¹¹

Finally, characterizing the equilibrium of the model is much easier under the stronger assumption that $\epsilon = 1$, so that per capita consumption and offspring receive equal weight in the utility function. In the presentation of the model, we will make this assumption and indicate at the appropriate time what happens when $\epsilon \neq 1$.

4.2. Technology

The consumption-capital good in the economy, final output Y , is produced according to

$$Y_t = A_t^\sigma K_t^\alpha L_{Yt}^{1-\alpha}, \quad (12)$$

where A is the stock of ideas in the economy, K is capital, L_Y is labor, and the parameters satisfy $\sigma > 0$ and $0 < \alpha < 1$. While this kind of production function is commonly used in economics, it incorporates a fundamental insight into the process of economic growth. Specifically, the production function exhibits increasing returns to scale because of the nonrivalry of ideas. The strength of increasing returns is measured by σ .

The technology for producing offspring has already been discussed. It turns out to be convenient to invert this production function in the analysis that follows. Individuals are endowed with one unit of labor, and generating a net fertility rate of $n \equiv \tilde{n} - d$ requires $l = \beta(n)$ units of time, where

11. This restriction is closely related to the restriction in dynamic general equilibrium business cycle models that consumption must enter in log form if consumption and leisure are additively separable (leisure per person does not need to enter in log form because it is not growing over time). Alternative approaches to fertility can relax this assumption.

$$\beta(n) \equiv \left(\frac{n+d}{b} \right)^{1/\psi}. \quad (13)$$

The time that individuals have left over to supply to the labor market is therefore $1 - \beta(n)$. Note that $\beta(0) > 0$ (some time is required to maintain a constant level of population to compensate for the deaths at rate d), and $\beta(n)$ is a convex function. In the optimization problem, we have individuals choose n rather than \tilde{n} , but of course the two choices are equivalent. With N identical agents in the economy, the total change in population in an economy with net fertility n is given by

$$\dot{N}_t = n_t N_t. \quad (14)$$

We start the economy at time 0 with $N_0 > 0$ given.

Capital accumulates in this economy in the form of assets owned by members of the dynastic family. Letting v denote the per capita stock of assets ($K = Nv$ is imposed later, and $K_0 > 0$ is assumed),

$$\dot{v}_t = (r_t - n_t)v_t + w_t [1 - \beta(n_t)] - c_t - f_t, \quad (15)$$

where r is the market return on assets, w is the wage rate per unit of labor, and f represents per capita lump-sum taxes collected by the government ($f \equiv F/N$).

The final component of the technology of the economy is the production of ideas. New ideas are produced by researchers according to

$$\dot{A}_t = \bar{\delta}_t L_{At}, \quad (16)$$

where L_A denotes labor engaged in research, and \dot{A} represents the measure of new ideas created at a point in time. The resource constraint on labor is

$$L_{At} + L_{Yt} = [1 - \beta(n_t)] N_t \equiv L_t. \quad (17)$$

While individual researchers, who are small relative to the total number of researchers, take $\bar{\delta}$ as given, it may, in fact, depend on features of the aggregate economy. The true relationship between new ideas and research is assumed to be given by

$$\dot{A}_t = \delta L_{At}^\lambda A_t^\phi, \quad (18)$$

where $\delta > 0$, $0 < \lambda \leq 1$, and $\phi < 1$ are parameters. This formulation allows for both positive and negative externalities in research. At a point in time, congestion or duplication in research may reduce the social value of a marginal unit of research, associated with $\lambda < 1$. In addition, the productivity of research today may depend either positively (knowledge spillovers) or negatively (fishing out) on the stock of ideas discovered in the past. Equation (18) therefore allows for increasing, constant, or decreasing returns to scale in the production of new ideas.

4.3. Market Structure

Romer (1990) and others have emphasized that ideas are nonrivalrous but partially excludable. The assumption that ideas are at least partially excludable allows inventors to capture some of the social value that they create. This feature, together with the increasing returns to scale implied by nonrivalry, leads Romer, Grossman and Helpman, and Aghion and Howitt to favor models with profit-maximizing entrepreneurs and imperfect competition—what we might call “Silicon Valley” models.

Here, we will make an alternative assumption that will have the flavor of growth through basic science. In particular, we assume that ideas are nonrivalrous and nonexcludable; that is, they are pure public goods.¹² This means that inventors cannot use the market mechanism to capture any of the social value they create. In the absence of some nonmarket intervention, no one would become a researcher because of the fundamental ineffectiveness of property rights over basic science, and there would be no growth.

This alternative assumption serves two purposes. The primary purpose is that it greatly simplifies the analysis of the decentralized model. We assume that all markets are perfectly competitive and then introduce a government to collect lump-sum taxes and use the revenues to fund research publicly. Of course, this case may also be of independent interest. Previous Silicon Valley-style models have analyzed the case in which research is undertaken by private entrepreneurs who are compensated through imperfectly competitive markets. This chapter explores the alternative extreme in which growth is associated with basic science undertaken by publicly funded scientists.

The government collects lump-sum taxes F from individuals and uses this revenue to hire research scientists at the market wage w . We assume that the government collects as much revenue as needed so that a constant fraction of the labor force, $0 < \bar{s} < 1$, is hired as researchers: that is, $L_A = \bar{s}L$.

4.4. Equilibrium

A competitive equilibrium in this model is a sequence of quantities $\{c_t, Y_t, K_t, A_t, v_t, L_{Yt}, L_{At}, N_t, n_t\}$, prices $\{w_t, r_t\}$, and lump-sum taxes $\{F_t\}$ such that:

1. The head of the dynastic family chooses $\{c_t, n_t\}$ to maximize dynastic utility in equation (10) subject to the laws of motion for asset accumulation (15) and population (14), taking $\{r_t, w_t, F_t\}$ and v_0 and N_0 as given.
2. Firms producing output rent capital K_t and labor L_{Yt} to maximize profits, taking the rental prices r_t and w_t and the stock of ideas A_t as given.
3. Markets clear at the prices $\{w_t, r_t\}$ and the taxes $\{F_t\}$. In particular, the stock of assets held by consumers V_t is equal to the total capital stock K_t , and the number of researchers is a constant fraction \bar{s} of the labor force.

12. See Shell (1966) for an early application of this approach.

We now characterize the competitive equilibrium in steady state, that is, when all variables are growing at constant (exponential) rates.

The first-order conditions from the utility maximization problem for individuals imply that the steady-state fertility rate chosen by the dynastic family satisfies¹³:

$$\frac{(r - g_Y)[v_t + w\beta'(n)]}{\tilde{N}_t} = \frac{u_{\tilde{N}_t}}{u_{c_t}}. \quad (19)$$

This equation is the dynamic equivalent of the condition that the marginal rate of transformation (the left-hand side) equals the marginal rate of substitution (the right-hand side) between people and consumption. The marginal rate of transformation is based on the cost to the individual of increasing fertility, which involves two terms. First, there is a capital-narrowing effect: adding to the population dilutes the stock of assets per person. Second, there is the direct cost of wages that are foregone in order to increase the population growth rate. The total cost is scaled by the size of the population so that it is measured in terms of bodies rather than as a rate of growth, and it is multiplied by the effective discount rate $r - g_Y$ to put it on a flow basis. This marginal rate of transformation is equal to the static marginal rate of substitution $u_{\tilde{N}}/u_c$ along the optimal balanced growth path.

This relationship makes it clear why a unit elasticity of substitution between people and consumption is required. The marginal rate of transformation on the left-hand side of equation (19) will end up being proportional to y/\tilde{N} , where y is per capita output Y/N . Therefore, the marginal rate of substitution must be proportional to c/\tilde{N} for a balanced growth path to exist; otherwise, the cost and the benefit of fertility will grow at different rates and the economy will be pushed to a corner. The equation also makes clear why the curvature $\beta''(n) > 0$ is required: with $\beta(n) = 1 - \beta n$, for example, equation (19) does not depend directly on n , and households will move to a corner solution.

Other first-order conditions characterizing the equilibrium are more familiar. For example, consumption growth satisfies the following Euler equation:

$$\dot{c}_t/c_t = r_t - n_t - \rho. \quad (20)$$

Also, the first-order conditions from the firm's profit-maximization problem, assuming no depreciation, are

$$r_t = \alpha Y_t/K_t$$

and

$$w_t = \frac{(1 - \alpha)Y_t}{L_{Y_t}} = (1 - \alpha)y_t \cdot \frac{1}{1 - \beta(n_t)} \cdot \frac{1}{1 - \bar{s}}. \quad (21)$$

13. Robert Barro, in work in progress, shows that with an intertemporal elasticity of substitution equal to one, if we replace $r - g_Y$ with ρ , this condition holds at all points in time, not just along a balanced growth path.

With these first-order conditions in mind, we are ready to characterize the steady-state growth rate of the economy. Along the balanced growth path, the key growth rates of the model are all given by the growth rate of the stock of ideas:

$$g_y = g_k = g_c = \frac{\sigma}{1 - \alpha} g_A, \quad (22)$$

where g_z denotes growth rate of some variable z along the balanced growth path, y is per capita income Y/N , and k is capital per person K/N .¹⁴

The growth rate of ideas, g_A , is found by dividing both sides of equation (18) by A :

$$\frac{\dot{A}_t}{A_t} = \delta \frac{L_{A_t}^\lambda}{A_t^{1-\phi}}.$$

Along a balanced growth path, the numerator and the denominator of the right-hand side of this expression must grow at the same rate, and this requirement pins down the growth rate of A as

$$g_A = \frac{\lambda}{1 - \phi} g_{L_A}.$$

Finally, along a balanced growth path, L_A must grow at the rate of growth of the population. Therefore,

$$g_A = \frac{\lambda n}{1 - \phi}. \quad (23)$$

Combining this result with equation (22), we see that

$$g_y = \gamma n, \quad (24)$$

where $\gamma \equiv [\sigma/(1 - \alpha)][\lambda/(1 - \phi)]$.

As in Jones (1995), the per capita growth rate of the economy is proportional to the population growth rate. This is a direct consequence of increasing returns to scale: With $\sigma = 0$, there is no per capita growth in the long run. Note that balanced growth in the presence of population growth in this model requires $\alpha < 1$ and $\phi < 1$. That is, the capital accumulation equation and the law of motion for ideas must both be less than linear in their own state variables; otherwise, growth explodes and the level of consumption and income is infinite in a finite amount of time.

14. This relationship is derived as follows. First, the constancy-of-consumption growth requires a constant interest rate and therefore a constant capital-output ratio, yielding the first equality. Second, the asset accumulation equation in (15) is simply a standard capital accumulation equation. For the capital stock to grow at a constant rate, the capital-consumption ratio must be constant, yielding the second equality. Finally, log differentiating the production function in (12) yields the last equality.




Figure 1. Solving for n^{DC} .

4.5. Fertility in the Decentralized Economy

The rate of population growth is determined by consumer optimization, as in equation (19). Using the fact that $\epsilon = 1$ and $r - g_Y = \rho$ along a balanced growth path, and substituting for the wage from equation (21), equation (19) can be written as

$$k_t + (1 - \alpha)y_t \cdot \frac{\beta'(n)}{1 - \beta(n)} \cdot \frac{1}{1 - \bar{s}} = \frac{1}{\rho} c_t. \quad (25)$$

Some algebra then shows that along a balanced growth path, the rate of fertility satisfies¹⁵

$$\frac{\beta'(n^{DC})}{1 - \beta(n^{DC})} = \frac{1 - \bar{s}}{\rho}. \quad (26)$$

The solution to this equation exists and is unique under the assumption that $\beta'(0) < (1 - \bar{s})/\rho$, as shown in Figure 1. Recall that the relationship in equation (24) that $g_Y = \gamma n$ then determines the growth rate of the economy along the balanced growth path.

The steady-state growth rate of the economy is directly proportional to the net fertility rate. This rate is smaller the higher is the rate of time preference ρ or the

15. Specifically, divide both sides of the equation by k and use the fact that $y/k = r/\alpha$ and $c/k = y/k - g_Y = (1 - \alpha)/\alpha * r + \rho$ along a balanced growth path.

higher is the cost of fertility $\beta(\cdot)$. Interestingly, the growth rate of the economy is *decreasing* rather than increasing in the fraction of the labor force devoted to research. This is very different from the results in previous idea-based growth models and reflects the fact that growth is driven by a different mechanism. Here, changes in research intensity affect long-run growth only through their effect on fertility. A larger research sector takes labor away from the alternative use of producing offspring, which reduces population growth and therefore reduces steady-state per capita growth. It is important to note that this long-run effect is quite different from the short-run effect. In the short run, an increase in the fraction of the labor force devoted to research will lead to more new ideas and a faster rate of growth. Only in the long run is the fertility effect apparent.

5. WELFARE AND A PLANNER PROBLEM

With more than one generation of agents, it is not obvious how to define social welfare: It depends on how one weights the utility of different generations. We focus on a narrower question: Does the allocation of resources achieved in the market economy maximize the utility of each dynastic family given the initial conditions that constrain their choices?

To maximize the welfare of a representative generation (the generation alive at time zero here), the social planner solves

$$\max_{\{c_t, s_t, n_t\}} U_0 = \int_0^{\infty} e^{-\rho t} u(c_t, \tilde{N}_{0,t}) dt, \quad (27)$$

subject to

$$\dot{k}_t = A_t^\sigma k_t^\alpha (1 - s_t)^{1-\alpha} [1 - \beta(n_t)]^{1-\alpha} - c_t - n_t k_t, \quad (28)$$

$$\dot{A}_t = \delta s_t^\lambda [1 - \beta(n_t)]^\lambda N_t^\lambda A_t^\phi, \quad (29)$$

and

$$\dot{N}_t = n_t N_t. \quad (30)$$

The first-order conditions from this maximization problem can be combined to yield several equations of interest. First, optimal consumption satisfies a standard Euler equation

$$\frac{\dot{c}_t}{c_t} = \alpha \frac{y_t}{k_t} - n_t - \rho. \quad (31)$$

Second, the first-order conditions together with the equations governing the law of motion for capital and ideas can be solved to yield optimal research intensity in the steady state:

$$s^{SP} = \frac{1}{1 + \psi^{SP}}, \quad (32)$$

where

$$\psi^{SP} = \frac{1 - \alpha}{\lambda \sigma} \left[\frac{\rho(1 - \phi)}{\lambda n} + 1 - \phi \right].$$

To solve for the steady-state rate of population growth, we follow the steps used for the decentralized model. The first-order conditions from the planner's problem can be combined to yield a condition analogous to equation (19) in steady state:

$$\left[k_t + (1 - \alpha)y_t \cdot \frac{\beta'(n^{SP})}{1 - \beta(n^{SP})} \cdot \frac{1}{1 - s^{SP}} \right] \frac{\rho}{\tilde{N}_t} = \frac{u_{\tilde{N}_t}}{u_{ct}} + \mu_{2t} \lambda \frac{\dot{A}_t}{\tilde{N}_t}, \quad (33)$$

where μ_2 is the shadow value of an idea [the co-state variable corresponding to equation (29)].

The distortion that affects fertility choice can be seen by comparing this equation to the corresponding condition in the decentralized model, either equation (19) or (25). Individual agents ignore the extra benefit associated with increasing returns to scale provided by additional population. This distortion is reflected by the presence of the second term on the right-hand side of equation (33), which corresponds to the utility value of the extra ideas created by an additional person.

Some additional algebra reveals that, along the balanced growth path, the optimal fertility rate satisfies¹⁶

$$\frac{\beta'(n^{SP})}{1 - \beta(n^{SP})} = \frac{1}{\rho}. \quad (34)$$

Finally, the optimal steady-state growth rate of per capita income is given by $g_y^{SP} = \gamma n^{SP}$.

A comparison of equations (26) and (34) indicates that steady-state fertility and growth are inefficiently too slow in the decentralized economy, as shown in Figure 2. This results from the fact that, as noted above, individuals ignore the economy-wide benefit of fertility that is associated with increasing returns to scale: A larger population generates more ideas that benefit all agents in the economy. This is the "Mozart effect" mentioned by Phelps (1968).

In more general models that I have explored, this result can be overturned. For example, when the kernel of the utility function is generalized to place a higher weight on offspring, that is, when $\epsilon > 1$, it is possible for the decentralized economy to have a fertility rate and therefore a growth rate that is inefficiently too high. This occurs if \bar{s} is sufficiently smaller than s^{SP} .¹⁷ Second, fertility and

16. Once again, divide both sides of the equation by k and use the fact that $y/k = r/\alpha$ and $c/k = y/k - g_y = (1 - \alpha)/\alpha * r + \rho$ along a balanced growth path.

17. To see part of the intuition, recall that from the standpoint of the decentralized economy, a lower research intensity increases fertility.




Figure 2. Comparing n^{DC} and n^{SP} .

growth can be inefficiently high in a model in which the Romer (1990) market structure is used instead of the perfectly competitive/basic science market structure. With the imperfectly competitive market structure of Romer, capital is underpaid relative to its marginal product so that some resources are available to compensate entrepreneurs. However, recall that part of the opportunity cost of fertility is the additional capital that must be provided to offspring. Imperfect competition reduces this cost and can lead to inefficiently high fertility and growth.

6. DISCUSSION

This extended example contains a number of predictions about long-run growth, some of which are found in earlier papers and some of which are new. First, the “scale effects” prediction that has been a key problem in many endogenous growth models turns out to be a key feature in this model. Increasing returns to scale implies that the scale of the economy will matter. Instead of affecting (counterfactually) the long-run growth rate, however, scale affects the long-run *level* of per capita income. Large populations generate more ideas than small populations, and because ideas are nonrivalrous, the larger number of ideas translates into higher per capita income. Endogenous growth in the scale of the economy through fertility leads to endogenous growth in per capita income.

Changes in government policies can affect the long-run growth rate by affecting the rate of fertility. For example, suppose that for each child, parents have to pay a fraction of their wages in taxes. Such a tax will reduce fertility and therefore reduce per capita growth.

Other policies can also affect population growth and per capita growth in the model, but the effects are often counterintuitive on the surface. Specifically, the imposition of many taxes in the model will increase rather than decrease long-run growth (though once again, the short-run effects and the welfare effects can go in the opposite direction). For example, a tax on labor income creates a wedge between working and child-rearing, the untaxed activity, and will increase fertility and per capita growth. A tax on capital reduces the opportunity cost of fertility by reducing the capital stock and wages and therefore will also increase growth. Finally, as we have already seen, an increase in an existing government subsidy to research will reduce long-run growth in the model. Note that increasing the research subsidy may easily be welfare improving here, but not, as is often argued, because it increases the long-run growth rate.¹⁸ In general, these results emphasize the important point that long-run growth and welfare are different and may even respond to policy changes in opposite directions.

What policies should the government follow in this model to obtain the socially optimal allocation of resources? At least in steady state, the policy turns out to be very simple and conventional, contrary to the counterintuitive results just mentioned. Suppose the government taxes labor income at rate τ_L and uses the revenue to fund research, with no lump-sum rebates or taxes. In this case, it is easy to show that steady-state fertility achieves its socially optimal level. Moreover, the share of labor employed in research, \bar{s} , is equal to the tax rate τ_L . Therefore, by choosing a labor income tax rate of $\tau_L = s^{SP}$, the fraction of labor working in research as well as the steady-state fertility and growth rate match the social optimum.

A final issue worth considering is the plausibility of the way endogenous fertility is modeled. The assumption of a unit elasticity of substitution between consumption and offspring in the dynastic utility function is crucial for delivering sustained exponential growth in population, and therefore in per capita income. However, it is far from clear that future population growth will actually be sustained. For example, fertility rates throughout the world appear to be falling and demographic projections by the U.S. Bureau of the Census and the World Bank suggest that world population may stabilize at some point far into the future—maybe the twenty-third century (Doyle, 1997).

Jones (2001) examines a model with an elasticity of substitution greater than one in a study of growth over the very long run. In this case, the model generates a demographic transition similar to that observed in the advanced countries of the world and, at least for some parameterizations, suggests that population levels may stabilize. An important prediction of such a model is that exponential growth in per capita incomes would not be sustained. This does not mean that growth would necessarily cease, however. For example, a constant number of researchers could potentially generate a constant number of new ideas, leading to arithmetical rather than exponential growth.

18. In all of the examples in this paragraph, it is assumed that the tax revenue collected is rebated lump sum to the agents. The behavioral changes result from the substitution effects; without the lump-sum rebates, the income effect will neutralize the substitution effect.

In this sense, the framework does not necessarily suggest that sustained exponential growth must continue forever. As indicated earlier, linearity in the population equation does not guarantee growth; this depends on fertility behavior. Sustained growth seems to be a good description of the advanced economies for the last century or so. But if the model is correct, the future of per capita growth will hinge on the ability of the world economy to continue to devote more and more quality-adjusted resources to the production of new ideas.

7. CONCLUSION

Recent research has led to a large number of potential explanations for the engine of economic growth. Distinguishing among these explanations is important, both from a scientific standpoint and from a policy perspective. Some explanations suggest that increases in public investment in physical capital would be appropriate and others point to subsidies to private investment. Some suggest that imperfect competition and incentives for innovation are key and others stress the formation of human capital. Some suggest that growth rates may be much higher in the future, and others say that they will be much lower.

In order to generate sustained exponential growth like that observed in the United States for the last 125 years, models of growth require a differential equation that is linear, or at least very nearly so. Following the suggestion of Solow (1956), this chapter proposes that a successful theory of economic growth should provide an intuitive and compelling justification for this crucial assumption.

After proposing this standard to which our future models should aspire, the chapter attempts to make some progress. We begin by pointing out that, taking population growth as a given, it is possible to understand the exponential growth in per capita income without appealing to any additional linearity. Instead, the increasing returns to scale associated with the nonrivalry of ideas combined with the historical presence of population growth implies per capita growth.

The remainder of the chapter explores a model in which both population growth and per capita growth emerge endogenously. The crucial linearity appears in the law of motion for population, and the chapter argues this is a more natural location for linearity than other locations considered in existing growth models. Each family chooses a number of children to have, \tilde{n} . With N such agents in the economy, the net increase in population is given by $\dot{N} = nN$, where $n = \tilde{n} - d$. In other words, in deciding how many children to have, individuals choose the *proportional* rate of increase in the population. The linearity of the law of motion for population results from the biological fact of nature that people reproduce in proportion to their number. By itself, however, this linearity is not sufficient to generate per capita growth.

The second key ingredient of the model is increasing returns to scale. In line with the reasoning of Romer (1990) and others, increasing returns also seems to be a fact of nature. Ideas are a central feature of the world we live in. Ideas are nonrivalrous. Nonrivalry implies increasing returns to scale. This line of reasoning, rather than placing the key linearity in the equation of motion for technological progress, is the fundamental insight of the idea-based growth models, according to the view in

this chapter. Endogenous fertility and increasing returns, both motivated from first principles, are the key ingredients in an explanation of sustained and endogenous per capita growth.

REFERENCES

- Acemoglu, D. (2001), "Labor- and Capital-Augmenting Technical Change," mimeo, MIT.
- Aghion, P. and P. Howitt (1992), "A Model of Growth through Creative Destruction," *Econometrica* 60(2):323–51.
- Arrow, K. J. (1962), "The Economic Implications of Learning by Doing," *Review of Economic Studies* 29:153–73.
- Barro, R. J. (1974), "Are Government Bonds Net Wealth?," *Journal of Political Economy* 82:1095–117.
- Barro, R. J. and G. S. Becker (1989), "Fertility Choice in a Model of Economic Growth," *Econometrica* 57(2):481–501.
- Barro, R. J. and X. Sala-i-Martin (1995), *Economic Growth*, New York: McGraw-Hill.
- Becker, G. S. and R. J. Barro (1988), "A Reformulation of the Economic Theory of Fertility," *Quarterly Journal of Economics* 103(1):1–25.
- Ben-David, D. and D. H. Papell (1995), "The Great Wars, the Great Crash, and Steady-State Growth: Some New Evidence about an Old Stylized Fact," *Journal of Monetary Economics* 36:453–75.
- Blanchard, O. J. (1985), "Debts, Deficits, and Finite Horizons," *Journal of Political Economy* 93(2):223–47.
- Cannon, E. (1998), "Economies of Scale and Constant Returns to Capital: A Neglected Early Contribution to the Theory of Economic Growth," mimeo, University of Bristol.
- Dasgupta, P. S. (1969), "On the Concept of Optimum Population," *Review of Economic Studies* 36:295–318.
- Dinopoulos, E. and P. Thompson (1998), "Schumpeterian Growth without Scale Effects," *Journal of Economic Growth* 3(4):313–35.
- Doyle, R. (1997), "By the Numbers, Global Fertility and Population," *Scientific American* 276:26.
- Drandakis, E. M. and E. S. Phelps (1966), "A Model of Induced Invention, Growth, and Distribution," *Economic Journal* 76(304):823–40.
- Grossman, G. M. and E. Helpman (1989), "Product Development and International Trade," *Journal of Political Economy* 97(6):1261–83.
- (1991), *Innovation and Growth in the Global Economy*, Cambridge: MIT Press.
- Howitt, P. (1999), "Steady Endogenous Growth with Population and R&D Inputs Growing," *Journal of Political Economy* 107(4):715–30.
- Jones, C. I. (1995), "R&D-Based Models of Economic Growth," *Journal of Political Economy* 103(4):759–84.
- (1999), "Growth: With or Without Scale Effects?," *American Economic Association Papers and Proceedings* 89:139–44.
- (2001), "Was an Industrial Revolution Inevitable? Economic Growth over the Very Long Run," *Advances in Macroeconomics* 1(2): Article 1. <http://www.bepress.com/bejm-advances/vol1/iss2/art1>.
- (2002), "Sources of U.S. Economic Growth in a World of Ideas," *American Economic Review* 92(1):220–39.
- Jones, L. E. and R. Manuelli (1990), "A Convex Model of Economic Growth: Theory and Policy Implications," *Journal of Political Economy* 98:1008–38.
- Judd, K. L. (1985), "On the Performance of Patents," *Econometrica* 53(3):567–85.
- Kennedy, C. M. (1964), "Induced Bias in Innovation and the Theory of Distribution," *Economic Journal* 74(295):541–47.

- Kortum, S. S. (1997), "Research, Patenting, and Technological Change," *Econometrica* 65(6):1389–419.
- Kremer, M. (1993), "Population Growth and Technological Change: One Million B.C. to 1990," *Quarterly Journal of Economics* 108(4):681–716.
- Lucas, R. E. (1988), "On the Mechanics of Economic Development," *Journal of Monetary Economics* 22(1):3–42.
- Mulligan, C. B. and X. Sala-i-Martin (1993), "Transitional Dynamics in Two-Sector Models of Endogenous Growth," *Quarterly Journal of Economics* 108(3):739–74.
- Nordhaus, W. D. (1969), "An Economic Theory of Technological Change," *American Economic Association Papers and Proceedings* 59:18–28.
- Peretto, P. (1998), "Technological Change and Population Growth," *Journal of Economic Growth* 3(4):283–311.
- Phelps, E. S. (1966), "Models of Technical Progress and the Golden Rule of Research," *Review of Economic Studies* 33:133–45.
- (1968), "Population Increase," *Canadian Journal of Economics* 1(3):497–518.
- Pitchford, J. (1972), "Population and Optimal Growth," *Econometrica* 40(1):109–36.
- Raut, L. K. and T. N. Srinivasan (1994), "Dynamics of Endogenous Growth," *Economic Theory* 4:777–90.
- Razin, A. and U. Ben-Zion (1975), "An Intergenerational Model of Population Growth," *American Economic Review* 65(5):923–33.
- Rebelo, S. (1991), "Long-Run Policy Analysis and Long-Run Growth," *Journal of Political Economy* 99:500–21.
- Romer, P. M. (1987), "Crazy Explanations for the Productivity Slowdown," in S. Fischer, ed., *NBER Macroeconomics Annual*, Cambridge: MIT Press.
- (1990), "Endogenous Technological Change," *Journal of Political Economy* 98(5):S71–102.
- (1995), "Comment on a Paper by T. N. Srinivasan," in *Growth Theories in Light of the East Asian Experience*, Chicago: University of Chicago Press.
- Segerstrom, P. (1998), "Endogenous Growth Without Scale Effects," *American Economic Review* 88(5):1290–310.
- Shell, K. (1966), "Toward a Theory of Inventive Activity and Capital Accumulation," *American Economic Association Papers and Proceedings* 56:62–68.
- Sheshinski, E. (1967), "Optimal Accumulation with Learning by Doing," in K. Shell, ed., *Essays on the Theory of Economic Growth*, Cambridge: MIT Press.
- Simon, J. L. (1981), *The Ultimate Resource*, Princeton: Princeton University Press.
- Solow, R. M. (1956), "A Contribution to the Theory of Economic Growth," *Quarterly Journal of Economics* 70:65–94.
- (1994), "Perspectives on Growth Theory," *Journal of Economic Perspectives* 8(1):45–54.
- Stiglitz, J. E. (1990), "Comments: Some Retrospective Views on Growth Theory," in P. Diamond, ed., *Growth/Productivity/Unemployment: Essays to Celebrate Robert Solow's Birthday*, Cambridge: MIT Press, pp. 50–69.
- Swan, T. W. (1956), "Economic Growth and Capital Accumulation," *The Economic Record* 32:334–61.
- Young, A. (1998), "Growth without Scale Effects," *Journal of Political Economy* 106(1):41–63.