Increasing Returns and Long-Run Growth

Paul M. Romer

University of Rochester

This paper presents a fully specified model of long-run growth in which knowledge is assumed to be an input in production that has increasing marginal productivity. It is essentially a competitive equilibrium model with endogenous technological change. In contrast to models based on diminishing returns, growth rates can be increasing over time, the effects of small disturbances can be amplified by the actions of private agents, and large countries may always grow faster than small countries. Long-run evidence is offered in support of the empirical relevance of these possibilities.

I. Introduction

Because of its simplicity, the aggregate growth model analyzed by Ramsey (1928), Cass (1965), and Koopmans (1965) continues to form the basis for much of the intuition economists have about long-run growth. The rate of return on investment and the rate of growth of per capita output are expected to be decreasing functions of the level of the per capita capital stock. Over time, wage rates and capital-labor ratios across different countries are expected to converge. Consequently, initial conditions or current disturbances have no long-run effect on the level of output and consumption. For example, an exog-

This paper is based on work from my dissertation (Romer 1983). An earlier version of this paper circulated under the title "Externalities and Increasing Returns in Dynamic Competitive Analysis." At various stages I have benefited from comments by James J. Heckman, Charles M. Kahn, Robert G. King, Robert E. Lucas, Jr., Sergio Rebelo, Sherwin Rosen, José A. Scheinkman (the chairman of my thesis committee), and the referees. The usual disclaimer applies. I gratefully acknowledge the support of NSF grant no. SES-8320007 during the completion of this work.

[*Journal of Political Economy*, 1986, vol. 94, no. 5]

enous reduction in the stock of capital in a given country will cause prices for capital assets to increase and will therefore induce an offsetting increase in investment. In the absence of technological change, per capita output should converge to a steady-state value with no per capita growth. All these presumptions follow directly from the assumption of diminishing returns to per capita capital in the production of per capita output.

The model proposed here offers an alternative view of long-run prospects for growth. In a fully specified competitive equilibrium, per capita output can grow without bound, possibly at a rate that is monotonically increasing over time. The rate of investment and the rate of return on capital may increase rather than decrease with increases in the capital stock. The level of per capita output in different countries need not converge; growth may be persistently slower in less developed countries and may even fail to take place at all. These results do not depend on any kind of exogenously specified technical change or differences between countries. Preferences and the technology are stationary and identical. Even the size of the population can be held constant. What is crucial for all of these results is a departure from the usual assumption of diminishing returns.

While exogenous technological change is ruled out, the model here can be viewed as an equilibrium model of endogenous technological change in which long-run growth is driven primarily by the accumulation of knowledge by forward-looking, profit-maximizing agents. This focus on knowledge as the basic form of capital suggests natural changes in the formulation of the standard aggregate growth model. In contrast to physical capital that can be produced one for one from forgone output, new knowledge is assumed to be the product of a research technology that exhibits diminishing returns. That is, given the stock of knowledge at a point in time, doubling the inputs into research will not double the amount of new knowledge produced. In addition, investment in knowledge suggests a natural externality. The creation of new knowledge by one firm is assumed to have a positive external effect on the production possibilities of other firms because knowledge cannot be perfectly patented or kept secret. Most important, production of consumption goods as a function of the stock of knowledge and other inputs exhibits increasing returns; more precisely, knowledge may have an increasing marginal product. In contrast to models in which capital exhibits diminishing marginal productivity, knowledge will grow without bound. Even if all other inputs are held constant, it will not be optimal to stop at some steady state where knowledge is constant and no new research is undertaken.

These three elements—externalities, increasing returns in the production of output, and decreasing returns in the production of new knowledge—combine to produce a well-specified competitive equilibrium model of growth. Despite the presence of increasing returns, a competitive equilibrium with externalities will exist. This equilibrium is not Pareto optimal, but it is the outcome of a well-behaved positive model and is capable of explaining historical growth in the absence of government intervention. The presence of the externalities is essential for the existence of an equilibrium. Diminishing returns in the production of knowledge are required to ensure that consumption and utility do not grow too fast. But the key feature in the reversal of the standard results about growth is the assumption of increasing rather than decreasing marginal productivity of the intangible capital good knowledge.

The paper is organized as follows. Section II traces briefly the history of the idea that increasing returns are important to the explanation of long-run growth and describes some of the conceptual difficulties that impeded progress toward a formal model that relied on increasing returns. Section III presents empirical evidence in support of the model proposed here. Section IV presents a stripped-down, two-period version of the model that illustrates the tools that are used to analyze an equilibrium with externalities and increasing returns. Section V presents the analysis of the infinite-horizon, continuous-time version of the model, characterizing the social optimum and the competitive equilibrium, both with and without optimal taxes.

The primary motivation for the choice of continuous time and the restriction to a single state variable is the ease with which qualitative results can be derived using the geometry of the phase plane. In particular, once functional forms for production and preferences have been specified, useful qualitative information about the dynamics of the social optimum or the suboptimal competitive equilibrium can be extracted using simple algebra. Section VI presents several examples that illustrate the extent to which conventional presumptions about growth rates, asset prices, and cross-country comparisons may be reversed in this kind of economy.

II. Historical Origins and Relation to Earlier Work

The idea that increasing returns are central to the explanation of long-run growth is at least as old as Adam Smith's story of the pin factory. With the introduction by Alfred Marshall of the distinction between internal and external economies, it appeared that this explanation could be given a consistent, competitive equilibrium interpretation. The most prominent such attempt was made by Allyn Young in his 1928 presidential address to the Economics and Statistics section of the British Association for the Advancement of Science

(Young 1969). Subsequent economists (e.g., Hicks 1960; Kaldor 1981) have credited Young with a fundamental insight about growth, but because of the verbal nature of his argument and the difficulty of formulating explicit dynamic models, no formal model embodying that insight was developed.

Because of the technical difficulties presented by dynamic models, Marshall's concept of increasing returns that are external to a firm but internal to an industry was most widely used in static models, especially in the field of international trade. In the 1920s the logical consistency and relevance of these models began to be seriously challenged, in particular by Frank Knight, who had been a student of Young's at Cornell.[1] Subsequent work demonstrated that it is possible to construct consistent, general equilibrium models with perfect competition, increasing returns, and externalities (see, e.g., Chipman 1970). Yet Knight was at least partially correct in objecting that the concept of increasing returns that are external to the firm was vacuous, an "empty economic box" (Knight 1925). Following Smith, Marshall, and Young, most authors justified the existence of increasing returns on the basis of increasing specialization and the division of labor. It is now clear that these changes in the organization of production cannot be rigorously treated as technological externalities. Formally, increased specialization opens new markets and introduces new goods. All producers in the industry may benefit from the introduction of these goods, but they are goods, not technological externalities.[2]

Despite the objections raised by Knight, static models of increasing returns with externalities have been widely used in international trade. Typically, firm output is simply assumed to be increasing, or unit cost decreasing, in aggregate industry output. See Helpman (1984) for a recent survey. Renewed interest in dynamic models of growth driven by increasing returns was sparked in the 1960s following the publication of Arrow's (1962) paper on learning by doing. In his model, the productivity of a given firm is assumed to be an increasing function of cumulative aggregate investment for the industry. Avoiding the issues of specialization and the division of labor, Arrow argued that increasing returns arise because new knowledge is discovered as investment and production take place. The increasing returns were external to individual firms because such knowledge became publicly known.

To formalize his model, Arrow had to face two problems that arise

[1] For an account of the development of Young's ideas and of his correspondence with Knight, see Blitch (1983).

[2] For a treatment of increasing returns based on specialization, see Ethier (1982). Although the model there is essentially static, it demonstrates how specialization can be introduced in a differentiated products framework under imperfect competition.

in any optimizing model of growth in the presence of increasing returns. The first, familiar from static models, concerns the existence of a competitive equilibrium; as is now clear, if the increasing returns are external to the firm, an equilibrium can exist. The second problem, unique to dynamic optimizing models, concerns the existence of a social optimum and the finiteness of objective functions. In a standard optimizing growth model that maximizes a discounted sum or integral over an infinite horizon, the presence of increasing returns raises the possibility that feasible consumption paths may grow so fast that the objective function is not finite. An optimum can fail to exist even in the sense of an overtaking criterion. In the model of Arrow and its elaborations by Levhari (1966*a*, 1966*b*) and Sheshinski (1967), this difficulty is avoided by assuming that output as a function of capital and labor exhibits increasing returns to scale but that the marginal product of capital is diminishing given a fixed supply of labor. As a result, the rate of growth of output is limited by the rate of growth of the labor force. Interpreted as an aggregate model of growth (rather than as a model of a specific industry), this model leads to the empirically questionable implication that the rate of growth of per capita output is a monotonically increasing function of the rate of growth of the population. Like conventional models with diminishing returns, it predicts that the rate of growth in per capita consumption must go to zero in an economy with zero population growth.

The model proposed here departs from both the Ramsey-Cass-Koopmans model and the Arrow model by assuming that knowledge is a capital good with an increasing marginal product. Production of the consumption good is assumed to be globally convex, not concave, as a function of stock of knowledge when all other inputs are held constant. A finite-valued social optimum is guaranteed to exist because of diminishing returns in the research technology, which imply the existence of a maximum, technologically feasible rate of growth for knowledge. This is turn implies the existence of a maximum feasible rate of growth for per capita output. Over time, the rate of growth of output may be monotonically increasing, but it cannot exceed this upper bound.

Uzawa (1965) describes an optimizing growth model in which both intangible human capital and physical capital can be produced. In some respects, the human capital resembles knowledge as described in this paper, but Uzawa's model does not possess any form of increasing returns to scale. Instead, it considers a borderline case of constant returns to scale with linear production of human capital. In this case, unbounded growth is possible. Asymptotically, output and both types of capital grow at the same constant rate. Other optimizing models took the rate of technological change as exogenously given (e.g., Shell

1967*b*). Various descriptive models of growth with elements similar to those used here were also proposed during the 1960s (e.g., Phelps 1966; von Wiezsäcker 1966; Shell 1967*a*). Knowledge is accumulated by devoting resources to research. Production of consumption goods exhibits constant returns as a function of tangible inputs (e.g., physical capital and labor) and therefore exhibits increasing returns as a function of tangible and intangible inputs. Privately produced knowledge is in some cases assumed to be partially revealed to other agents in the economy. Because the descriptive models do not use explicit objective functions, questions of existence are generally avoided, and a full welfare analysis is not possible. Moreover, these models tend to be relatively restrictive, usually constructed so that the analysis could be carried out in terms of steady states and constant growth rate paths.

Continuous-time optimization problems with some form of increasing returns are studied in papers by Weitzman (1970), Dixit, Mirrlees, and Stern (1975), and Skiba (1978). Similar issues are considered for discrete-time models in Majumdar and Mitra (1982, 1983) and Dechert and Nishimura (1983). These papers differ from the model here primarily because they are not concerned with the existence of a competitive equilibrium. Moreover, in all these papers, the technical approach used to prove the existence of an optimum is different from that used here. They rely on either bounded instantaneous utility $U(c)$ or bounds on the degree of increasing returns in the problem; for example, the production function $f(k)$ is assumed to be such that $f(k)/k$ is bounded from above. The results here do not rely on either of these kinds of restrictions; in fact, one of the most interesting examples analyzed in Section VI violates both of these restrictions. Instead, the approach used here relies on the assumptions made concerning the research technology; the diminishing returns in research will limit the rate of growth of the state variable. A general proof that restrictions on the rate of growth of the state variable are sufficient to prove the existence of an optimum for a continuous-time maximization problem with nonconvexities is given in Romer (1986).

Because an equilibrium for the model proposed here is a competitive equilibrium with externalities, the analysis is formally similar to that used in dynamic models with more conventional kinds of externalities (e.g., Brock 1977; Hochman and Hochman 1980). It also has a close formal similarity to perfect-foresight Sidrauski models of money demand and inflation (Brock 1975) and to symmetric Nash equilibria for dynamic games (e.g., Hansen, Epple, and Roberds 1985). In each case, an equilibrium is calculated not by solving a social planning problem but rather by considering the maximization problem of an individual agent who takes as given the path of some endogenously determined aggregate variable. In the conventional analysis of externalities, the focus is generally on the social optimum and the set of taxes necessary to support it as a competitive equilibrium. While this question is addressed for this growth model, the discussion places more stress on the characterization of the competitive equilibrium without intervention since it is the most reasonable positive model of observed historical growth. One of the main contributions of this paper is to demonstrate how the analysis of this kind of suboptimal equilibrium can proceed using familiar tools like a phase plane even though the equations describing the equilibrium cannot be derived from any stationary maximization problem.

III. Motivation and Evidence

Because theories of long-run growth assume away any variation in output attributable to business cycles, it is difficult to judge the empirical success of these theories. Even if one could resolve the theoretical ambiguity about how to filter the cycles out of the data and to extract the component that growth theory seeks to explain, the longest available time series do not have enough observations to allow precise estimates of low-frequency components or long-run trends. When data aggregated into decades rather than years are used, the pattern of growth in the United States is quite variable and is apparently still influenced by cyclical movements in output (see fig. 1). Cross-country comparisons of growth rates are complicated by the difficulty of controlling for political and social variables that appear to strongly influence the growth process. With these qualifications in mind, it is useful to ask whether there is anything in the data that should cause economists to choose a model with diminishing returns, falling rates of growth, and convergence across countries rather than an alternative without these features.

Consider first the long-run trend in the growth rate of productivity or per capita gross domestic product (GDP). One revealing way to consider the long-run evidence is to distinguish at any point in time between the country that is the "leader," that is, that has the highest level of productivity, and all other countries. Growth for a country that is not a leader will reflect at least in part the process of imitation and transmission of existing knowledge, whereas the growth rate of the leader gives some indication of growth at the frontier of knowledge. Using GDP per man-hour as his measure of productivity, Maddison (1982) identifies three countries that have been leaders since 1700, the Netherlands, the United Kingdom, and the United States. Table 1 reports his estimates of the rate of growth of productivity in each country during the interval when it was the leader. When the

TABLE 1

PRODUCTIVITY GROWTH RATES FOR LEADING COUNTRIES

Lead Country	Interval	Annual Average Compound Growth Rate of GDP per Man-Hour (%)
Netherlands	1700–1785	−.07
United Kingdom	1785–1820	.5
United Kingdom	1820–90	1.4
United States	1890–1979	2.3

SOURCE.—Maddison (1982).

productivity growth rate is measured over intervals several decades long and compared over almost 3 centuries, the evidence clearly suggests that it has been increasing, not decreasing. The rate of growth of productivity increases monotonically from essentially zero growth in eighteenth-century Netherlands to 2.3 percent per year since 1890 in the United States.

Similar evidence is apparent from data for individual countries over shorter horizons. Table 2 reports growth rates in per capita GDP for the United States over five subperiods from 1800 to 1978. (The raw data used here are from Maddison [1979].) These rates also suggest a positive rather than a negative trend, but measuring growth rates over 40-year intervals hides a substantial amount of year-to-year or even decade-to-decade variation in the rate of growth. Figure 1 presents the average growth rate over the interval 1800–1839 (for which no intervening data are available) and for the subsequent 14 decades. Identifying a long-run trend in rates measured over decades is more problematical in this case, but it is straightforward to apply a simple nonparametric test for trend.

Table 3 reports the results of this kind of test for trend in the per capita rate of growth in GDP for several countries using raw data

TABLE 2

PER CAPITA GROWTH IN THE UNITED STATES

Interval	Average Annual Compound Growth Rate of Real per Capita GDP (%)
1800–1840	.58
1840–80	1.44
1880–1920	1.78
1920–60	1.68
1960–78	2.47

SOURCE.—Raw data are from Maddison (1979).

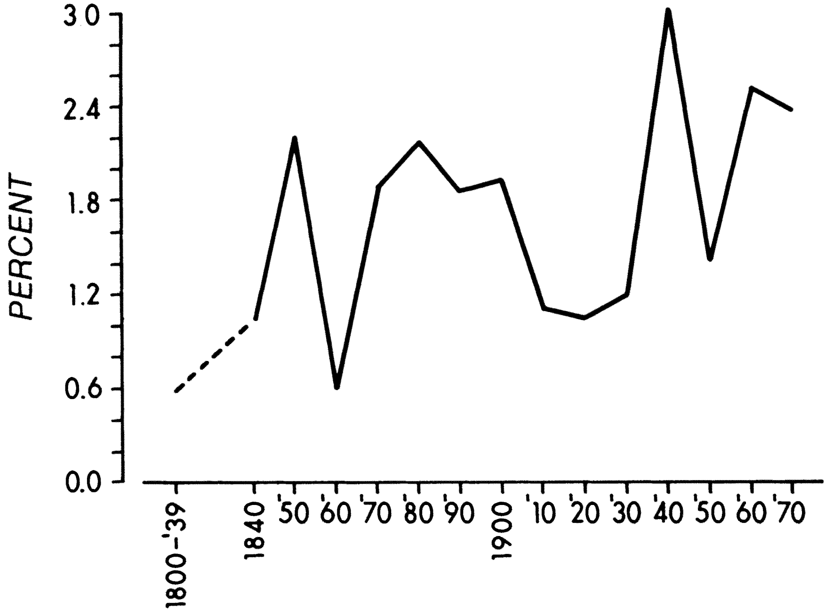

FIG. 1.—Average annual compound growth rate of per capita GDP in the United States for the interval 1800–1839 and for 14 subsequent decades. Data are taken from Maddison (1979).

from Maddison (1979). The sample includes all countries for which continuous observations on per capita GDP are available starting no later than 1870. As for the data for the United States graphed in figure 1, the growth rates used in the test for trend are measured over decades where possible. The statistic π gives the sample estimate of the probability that, for any two randomly chosen decades, the later decade has a higher growth rate.

Despite the variability evident from figure 1, the test for trend for the United States permits the rejection of the null hypothesis of a nonpositive trend at conventional significance levels. This is true even though growth over the 4 decades from 1800 to 1839 is treated as a single observation. However, rejection of the null hypothesis depends critically on the use of a sufficiently long data series. If we drop the observation on growth between 1800 and 1839, the estimate of π drops from .68 to .63 and the p-value increases from .03 to .11.[3] If we further restrict attention to the 11 decades from 1870 to 1978, π drops to .56 and the p-value increases to .29, so it is not surprising that studies that focus on the period since 1870 tend to emphasize the

[3] The p-value gives the probability of observing a value of π at least as large as the reported value under the null hypothesis that the true probability is .5.

TABLE 3

A TEST FOR TREND IN PER CAPITA GDP GROWTH RATES

	Date of First Observation	Number of Observations	π	p-Value
United Kingdom	1700	20	.63	.06
France	1700	18	.69	.01
Denmark	1818	16	.70	.02
United States	1800	15	.68	.03
Germany	1850	13	.67	.06
Sweden	1861	12	.58	.25
Italy	1861	12	.76	.01
Australia	1861	12	.64	.11
Norway	1865	12	.81	.002
Japan	1870	11	.67	.07
Canada	1870	11	.64	.12

NOTE.—π is the sample estimate for each country of the probability that, for any two growth rates, the later one is larger. The p-value is the probability of observing a value of π at least as large as the observed value under the null hypothesis that the true probability is .5. Except in the early years when data are sparse, per capita rates of growth of GDP were measured over successive decades. (Only two observations on growth rates are available for France prior to 1820; for the United Kingdom, only two prior to 1800; for the United States, only one from 1800 to 1840.) For the calculation of the p-value, see Kendall (1962). Data are from Maddison (1979).

constancy of growth rates in the United States. Rejection does not appear to depend on the use of the rate of growth in per capita GDP rather than the rate of growth of productivity. Reliable measures of the work force prior to 1840 are not available, but using data from Kuznets (1971) for the period 1840–1960 and from the 1984 Economic Report of the President for 1960–80, one can construct a similar test for trend in the rate of growth of productivity over successive decades. The results of this test, π equal to .64 with a p-value of .10, correspond closely to those noted above for growth in per capita GDP over the similar interval, 1840–1978.

Over the entire sample of 11 countries, the estimated value for π ranges from .58 to .81, with a p-value that ranges from .25 to .002. Five out of 11 of the p-values are less than .05, permitting rejection at the 5 percent level in a one-sided test of the null hypothesis that there is a nonpositive trend in the growth rate; eight out of 11 permit rejection at the 10 percent level.

For less developed countries, no comparable long-run statistics on per capita income are available. Reynolds (1983) gives an overview of the pattern of development in such countries. Given the paucity of precise data for less developed countries, he focuses on the "turning point" at which a country first begins to exhibit a persistent upward trend in per capita income. The timing of this transition and the pace of subsequent growth are strongly influenced by the variations in the world economy. A general pattern of historically unprecedented growth for the world economy is evident starting in the last part of the 1800s and continuing to the present. This general pattern is interrupted by a significant slowdown during the years between the two world wars and by a remarkable surge from roughly 1950 to 1973. Worldwide growth since 1973 has been slow only by comparison with that surge and appears to have returned to the high rates that prevailed in the period from the late 1800s to 1914.

Although all less developed countries are affected by the worldwide economy, the effects are not uniform. For our purposes, the key observation is that those countries with more extensive prior development appear to benefit more from periods of rapid worldwide growth and suffer less during any slowdown. That is, growth rates appear to be increasing not only as a function of calendar time but also as a function of the level of development. The observation that more developed countries appear to grow relatively faster extends to a comparison of industrialized versus less developed countries as well. In the period from 1950 to 1980, when official estimates for GDP are generally available, Reynolds reports that the median rate of growth of per capita income for his sample of 41 less developed countries was 2.3 percent, "clearly below the median for the OECD countries for the same period" (p. 975).

If it is true that growth rates are not negatively correlated with the level of per capita output or capital, then there should be no tendency for the dispersion in the (logarithm of the)[4] level of per capita income to decrease over time. There should be no tendency toward convergence. This contradicts a widespread impression that convergence in this sense has been evident, especially since the Second World War. Streissler (1979) offers evidence about the source of this impression and its robustness. For each year from 1950 to 1974, he measures the variance across countries of the logarithm of the level of per capita income. In a sample of ex post industrialized countries, those countries with a level of per capita income of at least $2,700 in 1974, clear evidence of a decrease in the dispersion over time is apparent. In a sample of ex ante industrialized countries, countries with a per capita income of at least $350 in 1950, no evidence of a decrease in the variance is apparent. The first sample differs from the second because it includes Japan and excludes Argentina, Chile, Ireland, Puerto Rico, and Venezuela. As one would expect, truncating the sample at the end biases the trend toward decreasing dispersion (and

[4] Examining the dispersion in the logarithm of the level of per capita income, not dispersion in the level itself, is the correct way to test for convergence in the growth rates. If the rate of growth were constant across countries that start from different levels, the dispersion in the logarithm of the levels will stay constant, but dispersion in the levels will increase.

at the beginning toward increasing dispersion). When a sample of all possible countries is used, there is no evidence of a decrease in variance, but the interpretation of this result is complicated by the changing number of countries in the sample in each year due to data limitations.

Baumol (1985) reports similar results. When countries are grouped into industrialized, intermediate, centrally planned, and less developed economies, he argues that there is a tendency toward convergence in the level of productivity within groups, even though there is no tendency toward overall convergence. The tendency toward convergence is clear only in his group of industrialized economies, which corresponds closely to the sample of ex post industrialized countries considered by Streissler. In any case, he finds no obvious pattern in his entire sample of countries; if anything, there is a weak tendency toward divergence.[5]

The other kind of evidence that bears directly on the assumption of increasing returns in production comes from growth accounting exercises and the estimation of aggregate production functions. Economists believe that virtually all technical change is endogenous, the outcome of deliberate actions taken by economic agents. If so and if production exhibits constant returns to scale, one would expect to be able to account for the rate of growth of output in terms of the rates of growth of all inputs. The difficulty in implementing a direct test of this assertion lies in correctly measuring all the inputs to production, especially for intangible capital inputs such as knowledge. In a comprehensive attempt to account for the rates of growth in output in terms of rates of growth of all inputs, including human and nonhuman, tangible and intangible stocks of capital, Kendrick (1976) concluded that rates of growth of inputs are not sufficient to explain the rate of growth of output in the 40-year interval 1929–69. For various sectors and levels of aggregation, the rate of growth of output is 1.06–1.30 times the appropriate aggregate measure of the rate of growth for inputs. This kind of estimate is subject to substantial, unquantified uncertainty and cannot be taken as decisive support for the presence of increasing returns. But given the repeated failure of this kind of growth accounting exercise, there is no basis in the data for excluding the possibility that aggregate production functions are best described as exhibiting increasing returns.

[5] Baumol (1985) argues that the convergence he observes among the industrialized countries results from a transmission process for knowledge that takes place among the industrialized countries but does not extend to centrally planned or less developed countries. He would not agree that the apparent convergence is an artifact of an ex post choice of the industrialized countries. Since he does not treat this issue directly, it is difficult to resolve it from his data. He does admit that his groupings are "somewhat arbitrary."

IV. A Simple Two-Period Model

Even in the presence of increasing returns and externalities, calculating a social optimum is conceptually straightforward since it is equivalent to solving a maximization problem. Standard mathematical results can be used to show that a maximum exists and to characterize the solution by means of a set of necessary conditions. Despite the presence of global increasing returns, the model here does have a social optimum. The next section illustrates how it can be supported as a competitive equilibrium using a natural set of taxes and subsidies. This optimum is of theoretical and normative interest, but it cannot be a serious candidate for describing the observed long-run behavior of per capita output. To the extent that appropriate taxes and subsidies have been used at all, they are a quite recent phenomenon.

The model here also has an equilibrium in the absence of any governmental intervention. Much of the emphasis in what follows focuses on how to characterize the qualitative features of this suboptimal dynamic equilibrium. Although it is suboptimal, the competitive equilibrium does satisfy a constrained optimality criterion that can be used to simplify the analysis much as the study of the social optimization problem simplifies the analysis in standard growth models.

The use of a constrained or restricted optimization problem is not a new approach to the analysis of a suboptimal dynamic equilibrium. For example, it has been widely used in the perfect-foresight models of inflation. Nonetheless, it is useful to describe this method in some detail because previous applications do not highlight the generality of the approach and because the dynamic setting tends to obscure its basic simplicity. Hence, I start by calculating a competitive equilibrium for a greatly simplified version of the growth model.

Specifically, consider a discrete-time model of growth with two periods. Let each of S identical consumers have a twice continuously differentiable, strictly concave utility function $U(c_1, c_2)$, defined over consumption of a single output good in periods 1 and 2. Let each consumer be given an initial endowment of the output good in period 1. Suppose that production of consumption goods in period 2 is a function of the state of knowledge, denoted by k, and a set of additional factors such as physical capital, labor, and so forth, denoted by a vector $\mathbf{x}$.[6] To restrict attention to a choice problem that is essentially

[6] For most of the subsequent discussion, k will be treated as a stock of disembodied knowledge, i.e., knowledge in books. This is merely an expositional convenience and is not essential. For example, if one wants to assume that all knowledge is embodied in some kind of tangible capital such as conventional physical capital or human capital, k can be reinterpreted throughout as a composite good made up of both knowledge and the tangible capital good.

one-dimensional, assume that only the stock of knowledge can be augmented; the factors represented by $\mathbf{x}$ are available in fixed supply. To capture the basic idea that there is a trade-off between consumption today and knowledge that can be used to produce more consumption tomorrow, assume that there is a research technology that produces knowledge from forgone consumption in period 1. Because the economy here has only two periods, we need not be concerned with the problem that arises in an infinite-horizon model when consumption grows too fast and discounted utility goes to infinity. Thus we do not need diminishing returns in research to limit the rate of growth of knowledge, and we can choose a simple linear technology with units such that one unit of forgone consumption produces one unit of knowledge. A more realistic diminishing returns research technology is described in the infinite-horizon model presented in the next section.

Since newly produced private knowledge can be only partially kept secret and cannot be patented, we can represent the technology of firm i in terms of a twice continuously differentiable production function F that depends on the firm-specific inputs k_i and $\mathbf{x}_i$ and on the aggregate level of knowledge in the economy. If N is the number of firms, define this aggregate level of knowledge as $K = \Sigma_{i=1}^{N} k_i$.

The first major assumption on the production function $F(k_i, K, \mathbf{x}_i)$ is that, for any fixed value of K, F is concave as a function of k_i and $\mathbf{x}_i$. Without this assumption, a competitive equilibrium will not exist in general. Once concavity is granted, there is little loss of generality in assuming that F is homogeneous of degree one as a function of k_i and $\mathbf{x}_i$ when K is held constant; any concave function can be extended to be homogeneous of degree one by adding an additional factor to the vector $\mathbf{x}$ if necessary (Rockafellar 1970, p. 67). McKenzie (1959) refers to this additional factor as an entrepreneurial factor. It can be interpreted as an accounting device that transforms any profits into factor payments.

By the homogeneity of F in k_i and $\mathbf{x}_i$ and by the assumption that F is increasing in the aggregate stock of knowledge, K, it follows that F exhibits increasing returns to scale. For any $\psi > 1$,

$$F(\psi k_i, \psi K, \psi \mathbf{x}_i) > F(\psi k_i, K, \psi \mathbf{x}_i) = \psi F(k_i, K, \mathbf{x}_i).$$

The second major assumption strengthens this considerably. It requires that F exhibit global increasing marginal productivity of knowledge from a social point of view. That is, for any fixed $\mathbf{x}$, assume that $F(k, Nk, \mathbf{x})$, production per firm available to a dictator who can set economywide values for k, is convex in k, not concave. This strengthening of the assumption of increasing returns is what distinguishes the production function used here from the one used in the models of Arrow, Levhari, and Sheshinski.

The equilibrium for the two-period model is a standard competitive equilibrium with externalities. Each firm maximizes profits taking K, the aggregate level of knowledge, as given. Consumers supply part of their endowment of output goods and all the other factors $\mathbf{x}$ to firms in period 1. With the proceeds, they purchase output goods in period 2. Consumers and firms maximize taking prices as given. As usual, the assumption that agents treat prices and the aggregate level K as given could be rationalized in a model with a continuum of agents. Here, it is treated as the usual approximation for a large but finite number of agents. Because of the externality, all firms could benefit from a collusive agreement to invest more in research. Although this agreement would be Pareto-improving in this model, it cannot be supported for the same reasons that collusive agreements fail in models without externalities. Each firm would have an incentive to shirk, not investing its share of output in research. Even if all existing firms could be compelled to comply, for example, by an economywide merger, new entrants would still be able to free-ride and undermine the equilibrium.

Because of the assumed homogeneity of F with respect to factors that receive compensation, profits for firms will be zero and the scale and number of firms will be indeterminate. Consequently, we can simplify the notation by restricting attention to an equilibrium in which the number of firms, N, equals the number of consumers, S. Then per firm and per capita values coincide. Assuming that all firms operate at the same level of output, we can omit firm-specific subscripts.

Let $\bar{\mathbf{x}}$ denote the per capita (and per firm) endowment of the factors that cannot be augmented; let $\bar{e}$ denote the per capita endowment of the output good in period 1. To calculate an equilibrium, define a family of restricted maximization problems indexed by K:

$$
\begin{aligned}
P(K):\quad & \max_{k\in[0,\,\bar{e}]} && U(c_1, c_2) \\
& \text{subject to} && c_1 \le \bar{e} - k, \\
& && c_2 \le F(k, K, \mathbf{x}), \\
& && \mathbf{x} \le \bar{\mathbf{x}}.
\end{aligned}
$$

Since U is strictly concave and $F(k, K, \mathbf{x})$ is concave in k and $\mathbf{x}$ for each value of K, $P(K)$ will have a unique solution k for each value of K. (The solution for $\mathbf{x}$ is trivially $\bar{\mathbf{x}}$.) In general, the implied values for c_1, c_2, and k have no economic meaning. If K differs from Sk, then $F(k, K, \bar{\mathbf{x}})$ is not a feasible level of per capita consumption in period 2. Equilibrium requires that the aggregate level of knowledge that is achieved in the economy be consistent with the level that is assumed when firms make production decisions. If we define a function $\Gamma: \mathbb{R} \to \mathbb{R}$ that sends K into S times the value of k that achieves the maximum for the problem $P(K)$, this suggests fixed points of Γ as candidates for equilibria.

To see that any fixed point K^* of Γ can indeed be supported as a competitive equilibrium, observe that $P(K^*)$ is a concave maximization problem with solution $k^* = K^*/S$, $c_1^* = \bar{e} - k^*$, and $c_2^* = F(k^*, Sk^*, \bar{\mathbf{x}})$. Since it is concave, standard necessary conditions for concave problems apply. Let $\mathcal{L}$ denote a Lagrangian for $P(K^*)$ with multipliers p_1, p_2, and w:

$$\mathcal{L} = U(c_1, c_2) + p_1(\bar{e} - k - c_1) + p_2[F(k, K, \mathbf{x}) - c_2] + w(\bar{\mathbf{x}} - \mathbf{x}).$$

When an interior solution is assumed, familiar arguments show that $p_j = D_jU(c_1^*, c_2^*)$ for $j = 1, 2$, that $p_1 = p_2D_1F(k^*, Sk^*, \bar{\mathbf{x}})$, and that $w = p_2D_3F(k^*, Sk^*, \bar{\mathbf{x}})$.[7] As always, the shadow prices w and p_j can be interpreted as equilibrium prices. To see this, consider first the maximization problem of the firm: $\max_k p_2F(k, Sk^*, \mathbf{x}) - p_1k - w \cdot \mathbf{x}$. Since the firm takes both prices and the aggregate level Sk^* as given, a trivial application of the sufficient conditions for a concave maximization problem demonstrates that k^* and $\bar{\mathbf{x}}$ are optimal choices for the firm. By the homogeneity of F with respect to its first and third arguments, profits will be zero at these values. Consider next the problem of the consumer. Income to the consumer will be the value of the endowment, $I = p_1\bar{e} + w \cdot \bar{\mathbf{x}} = p_2F(k^*, Sk^*, \bar{\mathbf{x}}) + p_1(\bar{e} - k^*)$. (The second equality follows from the homogeneity of F in k and $\mathbf{x}$.) When the necessary conditions $p_j = D_jU(c_1^*, c_2^*)$ from the problem $P(K^*)$ are used, it follows immediately that c_1^* and c_2^* are solutions to the problem max $U(c_1, c_2)$ subject to the budget constraint $p_1c_1 + p_2c_2 \leq I$. Note that the marginal rate of substitution for consumers will equal the private marginal rate of transformation perceived by firms, $D_1U(c_1^*, c_2^*)/D_2U(c_1^*, c_2^*) = D_1F(k^*, Sk^*, \bar{\mathbf{x}})$. Because of the externality, this differs from the true marginal rate of transformation for the economy, $D_1F(k^*, Sk^*, \bar{\mathbf{x}}) + SD_2F(k^*, Sk^*, \bar{\mathbf{x}})$.

Arguments along these lines can be used quite generally to show that a fixed point of a mapping like Γ defined by a family of concave problems $P(K)$ can be supported as a competitive equilibrium with externalities. The necessary conditions from a version of the Kuhn-Tucker theorem generate shadow prices associated with any solution to $P(K)$. The sufficient conditions for the problems of the consumer and the firm can then be used to show that the quantities from the

[7] Here, D denotes a derivative, D_i the partial derivative with respect to the ith argument.

solution will be chosen in an equilibrium in which these prices are taken as given. Conversely, an argument similar to the usual proof of the Pareto optimality of competitive equilibrium can be used to show that any competitive equilibrium with externalities for this kind of economy will satisfy the restricted optimality condition implicit in the problem $P(K)$ (Romer 1983). That is, if K^* is an equilibrium value of aggregate knowledge, then K^*/S will solve the problem $P(K^*)$. Thus equilibria are equivalent to fixed points of the function Γ.

This allows an important simplification because it is straightforward to characterize fixed points of Γ in terms of the underlying functions U and F. Substituting the constraints from $P(K)$ into the objective and using the fact that $\mathbf{x}$ will be chosen to be $\bar{\mathbf{x}}$, define a new function $V(k, K) = U(\bar{e} - k, F(k, K, \bar{\mathbf{x}}))$. Because of the increasing marginal productivity of knowledge, V is not a concave function; but for any fixed K, it is concave in k. Then the optimal choice of k in any problem $P(K)$ is determined by the equation $D_1V(k, K) = 0$. Fixed points of Γ are then given by substituting Sk for K and solving $D_1V(k, Sk) = 0$. Given functional forms for U and F, this equation can immediately be written in explicit form. The analysis can therefore exploit a three-way equivalence between competitive equilibria with externalities, fixed points of Γ, and solutions to an explicit equation $D_1V(k, Sk) = 0$.

The key observation in this analysis is that equilibrium quantities can be characterized as the solution to a concave maximization problem. Then prices can be generated from shadow prices or multipliers for this problem. The complete statement of the problem must be sought simultaneously with its solution because the statement involves the equilibrium quantities. But since $P(K)$ is a family of concave problems, solving simultaneously for the statement of the problem and for its solution amounts to making a simple substitution in a first-order condition.

V. Infinite-Horizon Growth

A. Description of the Model

The analysis of the infinite-horizon growth model in continuous time proceeds exactly as in the two-period example above. Individual firms are assumed to have technologies that depend on a path $K(t)$, $t \geq 0$, for aggregate knowledge. For an arbitrary path K, we can consider an artificial planning problem $P_\infty(K)$ that maximizes the utility of a representative consumer subject to the technology implied by the path K. Assume that preferences over the single consumption good take the usual additively separable, discounted form, $\int_0^\infty U(c(t))e^{-\delta t}dt$, with $\delta >$

0. The function U is defined over the positive real numbers and can have $U(0)$ equal to a finite number or to $-\infty$, for example, when $U(c) = \ln(c)$. Following the notation from the last section, let $F(k(t), K(t), \mathbf{x}(t))$ denote the instantaneous rate of output for a firm as a function of firm-specific knowledge at time t, economywide aggregate knowledge at time t, and the level of all other inputs at t. As before, we will assume that all agents take prices as given and that firms take the aggregate path for knowledge as given.

Additional knowledge can be produced by forgoing current consumption, but the trade-off is no longer assumed to be one-for-one. By investing an amount I of forgone consumption in research, a firm with a current stock of private knowledge k induces a rate of growth $\dot{k} = G(I, k)$. The function G is assumed to be concave and homogeneous of degree one; the accumulation equation can therefore be rewritten in terms of proportional rates of growth, $\dot{k}/k = g(I/k)$, with $g(y) = G(y, 1)$. A crucial additional assumption is that g is bounded from above by a constant α. This imposes a strong form of diminishing returns in research. Given the private stock of knowledge, the marginal product of additional investment in research, Dg, falls so rapidly that g is bounded. An inessential but natural assumption is that g is bounded from below by the value $g(0) = 0$. Knowledge does not depreciate, so zero research implies zero change in k; moreover, existing knowledge cannot be converted back into consumption goods. As a normalization to fix the units of knowledge, we can specify that $Dg(0) = 1$; one unit of knowledge is the amount that would be produced by investing one unit of consumption goods at an arbitrarily slow rate.

Assume as before that factors other than knowledge are in fixed supply. This implies that physical capital, labor, and the size of the population are held constant. If labor were the only other factor in the model, exponential population growth could be allowed at the cost of additional notation; but as was emphasized in the discussion of previous models, a key distinguishing feature of this model is that population growth is not necessary for unbounded growth in per capita income. For simplicity it is left out. Allowing for accumulation of physical capital would be of more interest, but the presence of two state variables would preclude the simple geometric characterization of the dynamics that is possible in the case of one state variable. If knowledge and physical capital are assumed to be used in fixed proportions in production, the variable $k(t)$ can be interpreted as a composite capital good. (This is essentially the approach used by Arrow [1962] in the learning-by-doing model.) Given increasing marginal productivity of knowledge, increasing marginal productivity of a composite k would still be possible if the increasing marginal productivity of knowledge were sufficient to outweigh the decreasing marginal productivity associated with the physical capital.

Within the restrictions imposed by tractability and simplicity, the assumptions on the technology attempt to capture important features of actual technologies. As noted in Section II, estimated aggregate production functions do appear to exhibit some form of increasing returns to scale. Assuming that the increasing returns arise because of increasing marginal productivity of knowledge accords with the plausible conjecture that, even with fixed population and fixed physical capital, knowledge will never reach a level where its marginal product is so low that it is no longer worth the trouble it takes to do research. If the marginal product of knowledge were truly diminishing, this would imply that Newton, Darwin, and their contemporaries mined the richest veins of ideas and that scientists now must sift through the tailings and extract ideas from low-grade ore. That knowledge has an important public good characteristic is generally recognized.[8] That the production of new knowledge exhibits some form of diminishing marginal productivity at any point in time should not be controversial. For example, even though it may be possible to develop the knowledge needed to produce usable energy from nuclear fusion by devoting less than 1 percent of annual gross national product (GNP) to the research effort over a period of 20 years, it is likely that this knowledge could not be produced by next year regardless of the size of the current research effort.

B. *Existence and Characterization of a Social Optimum*

Before using necessary conditions to characterize the solutions to either the social optimization problem, denoted as PS_∞, or any of the artificial optimization problems $P_\infty(K)$, I must verify that these problems have solutions. First I state the problems precisely. Let k_0 denote the initial stock of knowledge per firm for the economy. As in the last section, I will always work with the same number of firms and consumers. Because the choice of $\mathbf{x} = \bar{\mathbf{x}}$ is trivial, I suppress this argument, writing $f(k, K) = F(k, K, \bar{\mathbf{x}})$. Also, let $\mathcal{F}(k) = f(k, Sk) = F(k, Sk, \bar{\mathbf{x}})$ denote the globally convex (per capita) production function that would be faced by a social planner. In all problems that follow, the constraint $\dot{k}(t) \geq 0$ for all $t \geq 0$ and the initial condition $k(0) = k_0$ will be understood:

[8] See, e.g., Bernstein and Nadiri (1983) for estimates from the chemical industry suggesting that spillover effects can be quite large.

$$PS_\infty: \quad \max \int_0^\infty U(c(t))e^{-\delta t}dt$$

$$\text{subject to } \frac{\dot{k}(t)}{k(t)} = g\left(\frac{\mathcal{F}(k(t)) - c(t)}{k(t)}\right);$$

$$P_\infty(K): \quad \max \int_0^\infty U(c(t))e^{-\delta t}dt$$

$$\text{subject to } \frac{\dot{k}(t)}{k(t)} = g\left(\frac{f(k(t), K(t)) - c(t)}{k(t)}\right).$$

Note that the only difference between these two problems lies in the specification of the production function. In the first case, it is convex and invariant over time. In the second, it is concave but depends on time through its dependence on the path $K(t)$. I can now state the theorem that guarantees the existence of solutions to each of these problems.

THEOREM 1. Assume that each of U, f, and g is a continuous real-valued function defined on a subset of the real line. Assume that U and g are concave. Suppose that $\mathcal{F}(k) = f(k, Sk)$ satisfies a bound $\mathcal{F}(k) \leq \mu + k^\rho$ and that $g(z)$ satisfies the bounds $0 \leq g(\mathbf{x}) \leq \alpha$ for real numbers μ, ρ, and α. Then if $\alpha\rho$ is less than the discount factor δ, PS_∞ has a finite-valued solution, and $P_\infty(K)$ has a finite-valued solution for any path $K(t)$ such that $K(t) \leq K(0)e^{\alpha t}$.

The proof, given in an appendix available on request, amounts to a check that the conditions of theorem 1 in Romer (1986) are satisfied. Note that if α is less than δ the inequality $\alpha\rho < \delta$ allows for $\rho > 1$. Thus the socially feasible production function $\mathcal{F}$ can be globally convex in k, with a marginal social product and an average social product of knowledge that increase without bound.

The analysis of the social planning problem PS_∞ in terms of a current-valued Hamiltonian and a phase plane follows along familiar lines (see, e.g., Arrow 1967; Cass and Shell 1976*a*, 1976*b*). Define $H(k, \lambda) = \max_c U(c) + \lambda\{kg([\mathcal{F}(k) - c]/k)\}$. For simplicity, assume that the functions U, f, and g are twice continuously differentiable. The first-order necessary conditions for a path $k(t)$ to be a maximum for PS_∞ are that there exists a path $\lambda(t)$ such that the system of first-order differential equations $\dot{k} = D_2H(k, \lambda)$ and $\dot{\lambda} = \delta\lambda - D_1H(k, \lambda)$ are satisfied and that the paths satisfy two boundary conditions: the initial condition on k and the transversality condition at infinity, $\lim_{t\to\infty} \lambda(t)k(t)e^{-\delta t} = 0$.[9]

[9] Proving the necessity of the transversality condition for a maximization problem that is not concave takes relatively sophisticated mathematical methods. Ekeland and

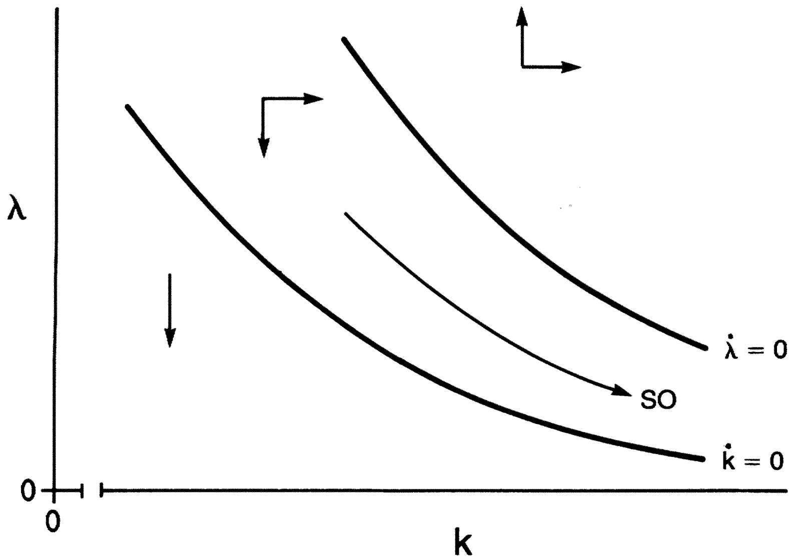

FIG. 2.—Geometry of the phase plane for a typical social optimum. Arrows indicate directions of trajectories in different sections of the plane. The rate of change of the stock of knowledge, $\dot{k}$, is zero everywhere on or below the locus denoted by $\dot{k} = 0$; *SO* denotes the socially optimal trajectory that stays everywhere between the lines $\dot{\lambda} = 0$ and $\dot{k} = 0$.

Under the assumption that $\lim_{c \to 0} DU(c) = \infty$, maximizing over c in the definition of $H(k, \lambda)$ implies that $DU(c) = \lambda Dg([\mathcal{F}(k) - c]/k)$ whenever the constraint $\dot{k} \geq 0$ is not binding; otherwise, $c = \mathcal{F}(k)$. This gives c as a function of k and λ. Substituting this expression in the equations for $\dot{k}$ and $\dot{\lambda}$ gives a system of first-order equations that depends only on k and λ.

Because of the restriction that $\dot{k}$ be nonnegative, the plane can be divided into two regions defined by $\dot{k} = 0$ and $\dot{k} \geq 0$ (see fig. 2). In a convenient abuse of the terminology, I will refer to the locus of points dividing these two regions as the $\dot{k} = 0$ locus. Along this locus, both the conditions $c = \mathcal{F}(k)$ and $DU(c) = \lambda Dg([\mathcal{F}(k) - c]/k)$ must hold. Thus the $\dot{k} = 0$ locus is defined by the equation $DU(\mathcal{F}(k)) = \lambda$. By the concavity of U, it must be a nonincreasing curve in the k-λ plane.

As usual, the equation $\dot{\lambda} = 0$ defines a simple locus in the plane. When the derivative $D_1H(k, \lambda)$ is evaluated along the $\dot{k} = 0$ locus, the equation for $\dot{\lambda}$ there can be written $\dot{\lambda}/\lambda = \delta - D\mathcal{F}(k)$. If $D\mathcal{F}$ increases without bound, there exists a value of $\hat{k}$ such that $D\mathcal{F}(k) > \delta$ for all k

Scheinkman (1983) prove the necessity of the transversality condition for nonconcave discrete-time problems. In continuous time, a proof that requires a local Lipschitz condition is given by Aubin and Clarke (1979).

larger than $\hat{k}$, and for all such k, the $\dot{\lambda} = 0$ locus lies above the $\dot{k} = 0$ locus. It may be either upward or downward sloping. If $\mathcal{F}$ were concave and satisfied the usual Inada conditions, $\dot{\lambda} = 0$ would cross $\dot{k} = 0$ from above and the resulting steady state would be stable in the usual saddle-point sense. Here, $\dot{\lambda} = 0$ may cross $\dot{k} = 0$ either from above or from below. If $D\mathcal{F}(k)$ is everywhere greater than δ, the $\dot{\lambda} = 0$ locus lies everywhere above the $\dot{k} = 0$ locus, and $\hat{k}$ can be taken to be zero. (This is the case illustrated in fig. 2.) Starting from any initial value greater than $\hat{k}$, the optimal trajectory $(\lambda(t), k(t))$, $t \geq 0$, must remain above the region where $\dot{k} = 0$. Any trajectory that crosses into this region can be shown to violate the transversality condition. Consequently, $k(t)$ grows without bound along the optimal trajectory.

This social optimum cannot be supported as a competitive equilibrium in the absence of government intervention. Any competitive firm that takes $K(t)$ as given and is faced with the social marginal products as competitive prices will choose not to remain at the optimal quantities even if it expects all other firms to do so. Each firm will face a private marginal product of knowledge (measured in terms of current output goods) equal to $D_1 f$; but the true shadow price of capital will be $D_1 f + SD_2 f > D_1 f$. Given this difference, each firm would choose to acquire less than the socially optimal amount of knowledge.

C. *Existence and Characterization of the Competitive Equilibrium*

Under a general set of conditions, this economy can be shown to have a suboptimal equilibrium in the absence of any intervention. It is completely analogous to the equilibrium for the two-period model. As in that model, it is straightforward to show that there is a three-way equivalence between competitive equilibria, fixed points of the mapping that sends a path $K(t)$ into S times the solution to $P_\infty(K)$, and solutions to an equation of the form $D_1 V(k, Sk) = 0$.[10] In the infinite-horizon case, this equation consists of a system of differential equations, which can be represented in terms of a phase plane, and a set of boundary conditions.

To derive these equations, consider the necessary conditions for the concave problem $P_\infty(K)$. Define a Hamiltonian, denoted as $\tilde{H}$ to distinguish it from the Hamiltonian H for the social planning problem PS_∞:

[10] An explicit proof of this result is given in Romer (1983). The method of proof is exactly as outlined in the two-period model. A generalized Kuhn-Tucker theorem is used to derive the necessary conditions that yield shadow prices for the maximization problems $P_\infty(K)$. Suppose K^* is a fixed point. If the consumer and the firm are faced with the shadow prices associated with $P_\infty(K^*)$, the sufficient conditions for their maximization problems are shown to be satisfied at the quantities that solve $P_\infty(K^*)$.

$$\tilde{H}(k, \lambda, K) = \max_c U(c) + \lambda\left[kg\left(\frac{f(k, K) - c}{k}\right)\right].$$

Then the necessary conditions for $k(t)$ to be a solution to $P_\infty(K)$ are that there exists a path $\lambda(t)$ such that $\dot{k}(t) = D_2\tilde{H}(k(t), \lambda(t), K(t))$ and $\dot{\lambda}(t) = \delta\lambda(t) - D_1\tilde{H}(k(t), \lambda(t), K(t))$ and such that the paths $k(t)$ and $\lambda(t)$ satisfy the boundary conditions $k(0) = k_0$ and $\lim_{t\to\infty} \lambda(t)k(t)e^{-\delta t} = 0$. Substituting $Sk(t)$ for $K(t)$ yields an autonomous system of differential equations, $\dot{k}(t) = D_2\tilde{H}(k(t), \lambda(t), Sk(t))$, $\dot{\lambda}(t) = \delta\lambda(t) - D_1\tilde{H}(k(t), \lambda(t), Sk(t))$, that can be characterized using the phase plane. The two boundary conditions must still hold. Any paths for $k(t)$ and $\lambda(t)$ that satisfy these equations and the boundary conditions will correspond to a competitive equilibrium, and all competitive equilibria can be characterized this way.

Before considering phase diagrams, I must show that a competitive equilibrium exists for some class of models. Standard results concerning the existence of solutions of differential equations can be used to prove that the equations for $\dot{\lambda}$ and $\dot{k}$ determine a unique trajectory through any point (k, λ) in the phase plane. The difficulty arises in showing that for any given value of k_0 there exists some value of λ_0 such that the transversality condition at infinity is satisfied along the trajectory through (k_0, λ_0). As opposed to the case in which these equations are generated by a concave maximization problem known to have a solution, there is no assurance that such a λ_0 exists.

The basic idea in the proof that such a λ_0 exists, and hence that a competitive equilibrium exists, is illustrated in example 1 from the next section. To state the general result, I need additional conditions that characterize the asymptotic behavior of the functions f and g. This is accomplished by means of an asymptotic exponent as defined by Brock and Gale (1969). Given a function $h(y)$, define the asymptotic exponent e of h as $e = \lim_{y\to\infty} \log_y|h(y)|$. Roughly speaking, $h(y)$ behaves asymptotically like the power function y^e. Also, recall that α is the maximal rate of growth for k implied by the research technology.

THEOREM 2. In addition to the assumptions of theorem 1, assume that U, f, and g are twice continuously differentiable. Assume also that $\mathcal{F}(k) = f(k, Sk)$ has an asymptotic exponent ρ such that $\rho > 1$ and $\alpha\rho < \delta$. Finally, assume that $Dg(\mathbf{x})$ has an asymptotic exponent strictly less than -1. Let $\tilde{k}$ be such that $D_1 f(k, Sk) > \delta$ for all $k > \tilde{k}$. Then if $k_0 > \tilde{k}$, there exists a competitive equilibrium with externalities in which $c(t)$ and $k(t)$ grow without bound.

The proof is given in Romer (1983, theorem 3). The assumption on the asymptotic growth of $\mathcal{F}$ is self-explanatory. The assumption on the asymptotic exponent of Dg is sufficient to ensure the boundedness of g. The condition on $D_1 f$ will be satisfied in most cases in which $\mathcal{F}(k)$

$= f(k, Sk)$ is convex. Examples of functions satisfying these assumptions are given in the next section.

Once the conditions for the existence of a competitive equilibrium have been established, the analysis reduces once again to the study of the phase plane summarizing the information in the differential equations. In many respects, this analysis is similar to that for the social optimum for this economy. The phase plane can once again be divided into regions where $\dot{k} = 0$ and $\dot{k} > 0$. Since by definition $\mathcal{F}(k) = f(k, Sk)$, the equations for c as a function of k and λ will be identical to those in the social optimum: $DU(c) = \lambda Dg([f(k, Sk) - c]/k)$ if $\dot{k} > 0$, $c = f(k, Sk)$ if $\dot{k} = 0$. As a result, the boundary locus for the region $\dot{k} = 0$ will also be identical with that from the social optimum. The only difference arises in the equation for $\dot{\lambda}$. Although the equality $H(k, \lambda) = \tilde{H}(k, \lambda, Sk)$ does hold, the derivatives $D_1H(k, \lambda)$ and $D_1\tilde{H}(k, \lambda, Sk)$ differ. In the first case, a term involving the expression $D\mathcal{F}(k) = D_1f(k, Sk) + SD_2f(k, Sk)$ will appear. In the second case, only the first part of this expression, $D_1f(k, Sk)$, appears. Therefore, $D_1H(k, \lambda)$ is always larger than $D_1\tilde{H}(k, \lambda, Sk)$. Consequently, the $\dot{\lambda} = 0$ locus for the competitive equilibrium must lie below that for the social optimum.

As was true of the social optimum, the $\dot{\lambda} = 0$ locus can be either upward or downward sloping. If $D_1f(k, Sk) > \delta$ for all k greater than some value $\tilde{k}$, the $\dot{\lambda} = 0$ locus will lie above $\dot{k} = 0$ for values of k to the right of $\tilde{k}$. Then the qualitative analysis is the same as that presented for the social optimum. Starting from an initial value $k_0 > \tilde{k}$, the only candidate paths for equilibria are ones that stay above the $\dot{k} = 0$ region; as before, paths that cross into this region will violate the transversality condition. A trajectory lying everywhere in the region where $\dot{k} > 0$ can fail to have $k(t)$ grow without bound only if the trajectory asymptotically approaches a critical point where $\dot{\lambda}$ and $\dot{k}$ are both zero, but no such point exists to the right of $\tilde{k}$. Hence, all the trajectories that are possible candidates for an equilibrium have paths for $k(t)$ that grow without bound. The existence result in theorem 2 shows that at least one such path satisfies the transversality condition at infinity.

D. *Welfare Analysis of the Competitive Equilibrium*

The welfare analysis of the competitive equilibrium is quite simple. The intuition from simple static models with externalities or from the two-period model presented in Section III carries over intact to the dynamic model here. In the calculation of the marginal productivity of knowledge, each firm recognizes the private return to knowledge, $D_1f(k, Sk)$, but neglects the effect due to the change in the aggregate level, $SD_2f(k, Sk)$; an increase in k induces a positive external effect

$D_2f(k, Sk)$ on each of the S firms in the economy. Consequently, the amount of consumption at any point in time is too high in the competitive equilibrium and the amount of research is too low. Any intervention that shifts the allocation of current goods away from consumption and toward research will be welfare-improving. As in any model with externalities, the government can achieve Pareto improvements not available to private agents because its powers of coercion can be used to overcome problems of shirking.

If the government has access to lump-sum taxation, any number of subsidy schemes will support the social optimum. Along the paths $k^*(t)$ and $\lambda^*(t)$ from the social optimum, taxes and subsidies must be chosen so that the first partial derivative of the Hamiltonian for the competitive equilibrium with taxes equals the first partial derivative of the Hamiltonian for the social planning problem; that is, the taxes and subsidies must be chosen so that the after-tax private marginal product of knowledge is equal to the social marginal product. This can be accomplished by subsidizing holdings of k, subsidizing accumulation $\dot{k}$, or subsidizing output and taxing factors of production other than k. The simplest scheme is for the government to pay a time-varying subsidy of $\sigma_1(t)$ units of consumption goods for each unit of knowledge held by the firm. If this subsidy is chosen to be equal to the term neglected by private agents, $\sigma_1(t) = SD_2f(k^*(t), Sk^*(t))$, private and social marginal products will be equal. A subsidy $\sigma_2(t)$ paid to a firm for each unit of goods invested in research would be easier to implement but is harder to characterize. In general, solving for $\sigma_2(t)$ requires the solution of a system of differential equations that depends on the path for $k^*(t)$. In the special case in which production takes the form $f(k, K) = k^{\nu}K^{\gamma}$, the optimal subsidy can be shown to be constant, $\sigma_2 = \gamma/(\nu + \gamma)$. (This calculation is also included in the app. available on request.)

While it is clear that the social marginal product of knowledge is greater than the private marginal product in the no-intervention competitive equilibrium, this does not necessarily imply that interest rates in the socially optimal competitive equilibrium with taxes will be higher than in the suboptimal equilibrium. In each case, the real interest rate on loans made in units of output goods can be written as $r(t) = -(\dot{p}/p)$, where $p(t) = e^{-\delta t}DU(c(t))$ is the present value price for consumption goods at date t. When utility takes the constant elasticity form $U(c) = [c^{(1-\theta)} - 1]/(1 - \theta)$, this reduces to $r(t) = \delta + \theta(\dot{c}/c)$. In the linear utility case in which $\theta = 0$, r will equal δ regardless of the path for consumption and in particular will be the same in the two equilibria. This can occur even though the marginal productivity of knowledge differs because the price of knowledge in terms of consumption goods (equal to the marginal rate of transformation between knowledge and consumption goods) can vary. Holders of knowledge earn capital gains and losses as well as a direct return equal to the private marginal productivity of knowledge. In the case of linear utility, these capital gains and losses adjust so that interest rates stay the same.

This logical point notwithstanding, it is likely that interest rates will be higher in the social optimum. On average, $\dot{c}/c$ will be higher in the social optimum; higher initial rates of investment with lower initial consumption must ultimately lead to higher levels of consumption. If there is any curvature in the utility function U, so that θ is positive, interest rates in the optimum will be greater than in the no-intervention equilibrium. In contrast to the usual presumption, cost-benefit calculations in a suboptimal equilibrium should use a social rate of discount that is higher than the market rate of interest.

VI. Examples

To illustrate the range of behavior possible in this kind of model, this section examines specific functional forms for the utility function U, the production function f, and the function g describing the research technology. Because the goal is to reach qualitative conclusions with a minimum of algebra, the choice of functional form will be guided primarily by analytical convenience. For the production function, assume that f takes the form noted above, $f(k, K) = k^{\nu}K^{\gamma}$. This is convenient because it implies that the ratio of the private and social marginal products,

$$\frac{D_1 f(k, Sk)}{D_1 f(k, Sk) + SD_2 f(k, Sk)} = \frac{\nu}{\nu + \gamma},$$

is constant. Nonincreasing private marginal productivity implies that $0 < \nu \leq 1$; increasing social marginal productivity implies that $1 < \gamma + \nu$. With these parameter values, this functional form is reasonable only for large values of k. For small values of k, the private and social marginal productivity of knowledge is implausibly small; at $k = 0$, they are both zero. This causes no problem provided we take a moderately large initial k_0 as given. An analysis starting from k_0 close to zero would have to use a more complicated (and more reasonable) functional form for f.

Recall that the rate of increase of the stock of knowledge is written in the homogeneous form $\dot{k} = G(I, k) = kg(I/k)$, where I is output minus consumption. The requirements on the concave function g are the normalization $Dg(0) = 1$ and the bound $g(I/k) < \alpha$ for all I/k. An analytically simple form satisfying these requirements is $g(z) = \alpha z/(\alpha + z)$. Recalling that δ is the discount rate, note that the bound required for the existence of a social optimum as given in theorem 1 requires the additional restriction that $\alpha(\nu + \gamma) < \delta$. Given the stated parameter restrictions, it is easy to verify that f and g satisfy all the requirements of theorems 1 and 2.

A. *Example 1*

With this specification of the technology for the economy, we can readily examine the qualitative behavior of the model for logarithmic utility $U(c) = \ln(c)$. The Hamiltonian can then be written as

$$\tilde{H}(k, \lambda, K, c) = \ln(c) + \lambda k g\left(\frac{f(k, K) - c}{k}\right).$$

Along (the boundary of the region in which) $\dot{k} = 0$, $Dg(0) = 1$ implies that $c = \lambda^{-1}$, so $\dot{k} = 0$ is determined by the equation

$$\lambda = [f(k, Sk)]^{-1} = S^{-\gamma}k^{-(\nu+\gamma)}.$$

The exact form for the locus $\dot{\lambda} = 0$ is algebraically complicated, but it is straightforward to show that, for large k, $\dot{\lambda} = 0$ lies above the $\dot{k} = 0$ locus since $D_1 f(k, Sk)$ will be greater than δ. Also, if we define the curve L_1 in the phase plane by the equation $\lambda = [1/(\delta - \alpha)]k^{-1}$, the $\dot{\lambda} = 0$ locus must cross L_1 from above as indicated in figure 3. (Details are given in the app. available on request.) Thus $\dot{k} = 0$ behaves as k to the power $-(\nu + \gamma) < -1$, and $\dot{\lambda} = 0$ is eventually trapped between $\dot{k} = 0$ and a line described by k to the power -1. In figure 3, representative trajectories t_1 and t_2 together with the competitive equilibrium trajectory CE are used to indicate the direction of trajectories in the various parts of the plane instead of the usual arrows.

Because the line L_1 is of the form $\lambda = [1/(\delta - \alpha)]k^{-1}$, any trajectory that eventually remains below L_1 will satisfy the transversality condition $\lim_{t\to\infty} e^{-\delta t}k(t)\lambda(t) = 0$. Given the geometry of the phase plane, it is clear that there must exist a trajectory that always remains between the loci $\dot{\lambda} = 0$ and $\dot{k} = 0$. Given the initial value k_0, index by the value of λ all the trajectories that start at a point (k_0, λ) between the two loci. The set of λ's corresponding to trajectories that cross $\dot{\lambda} = 0$ can have no smallest value, the set of λ's that correspond to trajectories that cross $\dot{k} = 0$ can have no largest value, and the two sets must be disjoint. Thus there exists a value λ_0 such that the trajectory through (k_0, λ_0) crosses neither locus and must therefore correspond to an equilibrium.[11]

[11] This is the essence of the proof of theorem 2.

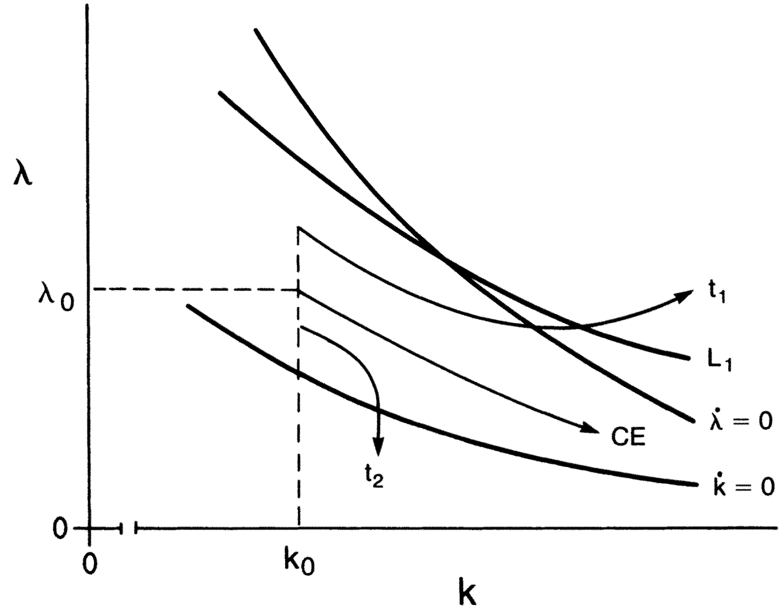

FIG. 3.—Geometry of the competitive equilibrium for example 1. The line L_1 is defined by the equation $\lambda = 1/(\delta - \alpha)k$; t_1 and t_2 denote representative trajectories in the phase plane; *CE* denotes the competitive equilibrium trajectory, which stays everywhere between the $\dot{\lambda} = 0$ and $\dot{k} = 0$ loci; λ_0 denotes the initial shadow price of knowledge corresponding to the initial stock of knowledge k_0.

In fact, the path resembles a conventional equilibrium in which the trajectory remains between the $\dot{\lambda} = 0$ and $\dot{k} = 0$ loci as it converges to a saddle point, although here it is as if the saddle point has been moved infinitely far to the right. Since the optimal trajectory cannot stop, capital grows without bound. Since the trajectory is downward sloping and since consumption is increasing in k and decreasing in λ, it is easy to see that consumption also grows without bound. Because of the difficulty of the algebra, it is not easy to describe the asymptotic rates of growth.

B. Example 2

Suppose now that utility is linear, $U(c) = c$. In the algebra and in the phase plane for this case, we can ignore the restriction $c \geq 0$ since it will not be binding in the region of interest. Maximizing out c from the Hamiltonian $\tilde{H}(k, \lambda, K, c) = c + \lambda k g((f - c)/k)$ implies that $c = f - \alpha k(\lambda^{.5} - 1)$. Then $f - c$ is positive (hence $\dot{k}$ is positive) if and only if $\lambda > 1$.

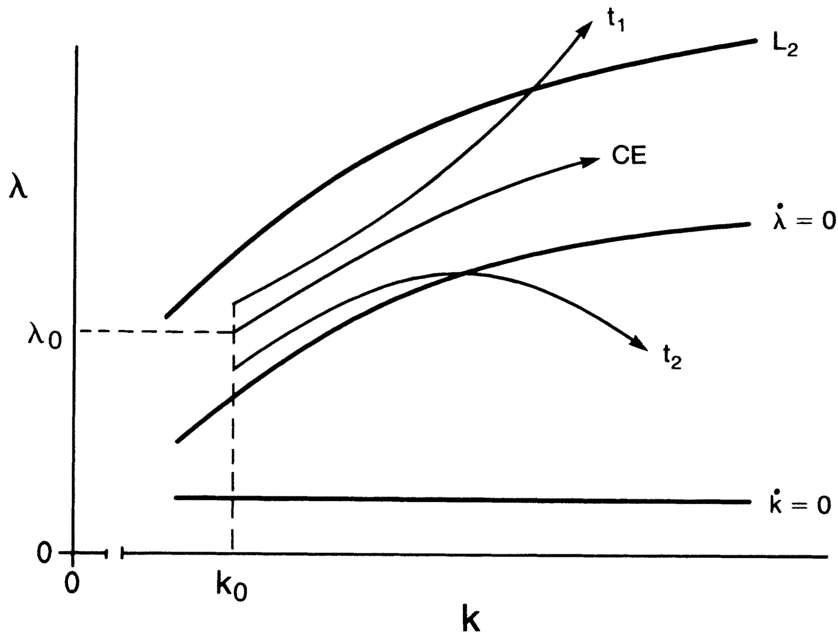

FIG. 4.—Geometry of the competitive equilibrium for example 2. The line L_2 is defined by an equation of the form $\lambda = Ak^{\nu+\gamma-1}$; t_1 and t_2 denote representative trajectories in the phase plane; CE denotes the competitive equilibrium trajectory that stays everywhere between L_2 and $\dot{\lambda} = 0$; λ_0 denotes the initial shadow price of knowledge.

In this example, it is possible to put tighter bounds on the behavior of the $\dot{\lambda} = 0$ locus and, more important, on the behavior of the equilibrium trajectory. As demonstrated in the appendix (available on request), $\dot{\lambda} = 0$ is upward sloping and behaves asymptotically like the power function $\lambda = Bk^{\nu+\gamma-1}$ for some constant B. For this economy, the equilibrium trajectory will lie above the $\dot{\lambda} = 0$ locus, so it is convenient to define an additional curve that will trap the equilibrium trajectory from above. For an appropriate choice of the constant A, the line L_2 defined by $\lambda = Ak^{\nu+\gamma-1}$ will lie above $\dot{\lambda} = 0$ and will have the property that trajectories must cross it from below (see fig. 4). Since trajectories must cross $\dot{\lambda} = 0$ from above, the same geometric argument as used in the last example demonstrates that there exists a trajectory that remains between these two lines. Consequently it must also behave asymptotically like $k^{\nu+\gamma-1}$. Since $k(t)$ can grow no faster than $e^{\alpha t}$, the product $\lambda(t)k(t)$ will be bounded along such a trajectory by a function of the form $e^{\alpha(\nu+\gamma)t}$. Since $\delta > (\nu + \gamma)\alpha$, this trajectory satisfies the transversality condition and corresponds to an equilibrium.

Along the equilibrium trajectory, λ behaves asymptotically like

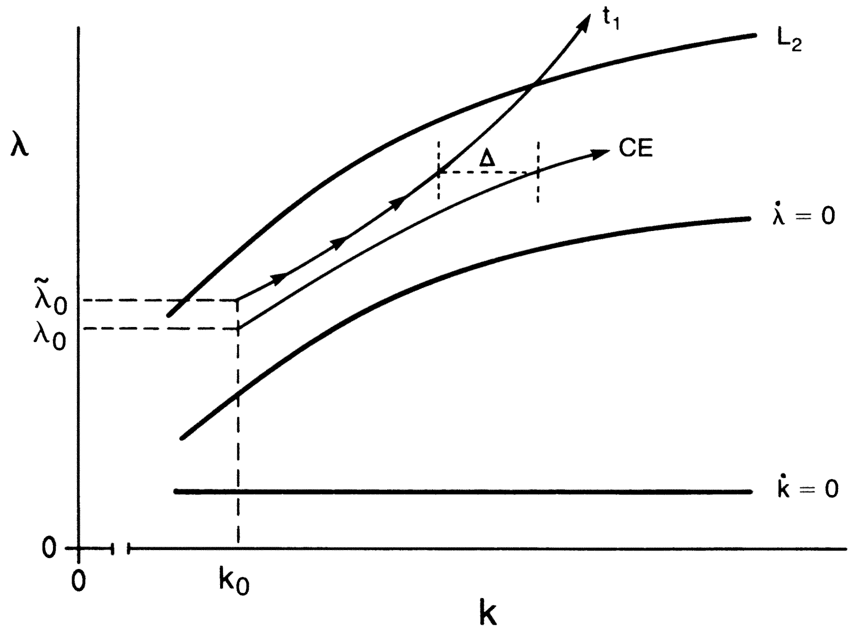

FIG. 5.—Geometry for the economy in example 2 when an exogenous increase of size Δ in the stock of knowledge is known to occur at a time $T > 0$. The equilibrium trajectory moves along t_1 until time T, at which point it is Δ units to the left of the trajectory CE. At time T, the economy jumps horizontally to CE with the change in the capital stock, but the path for $\lambda(t)$ is continuous. The equilibrium then proceeds along CE. $\tilde{\lambda}_0$ denotes the initial shadow price of knowledge in the case in which the exogenous increase will take place; λ_0 denotes the lower value that obtains in an economy in which no exogenous increase will take place.

$k^{\nu+\gamma-1}$. Given the expression noted above for c in terms of λ and k, c behaves asymptotically like $k^{\nu+\gamma} - \alpha k^{1+(.5)(\nu+\gamma-1)}$ and $I = f - c$ behaves like $k^{1+(.5)(\nu+\gamma-1)}$. Then c, I, C/k, and I/k go to infinity with k. By the assumptions on the research technology, I/k going to infinity implies that $\dot{k}/k$ approaches its upper bound α. Consequently, the percentage rate of growth of output and of consumption will be increasing, both approaching the asymptotic upper bound $\alpha(\nu + \gamma)$.

Because the equilibrium trajectory is upward sloping, this economy will exhibit different stability properties from either the conventional model or the economy with logarithmic utility described above. Figure 5 illustrates a standard exercise in which a perfect-foresight equilibrium is perturbed. Suppose that at time 0 it is known that the stock of knowledge will undergo an exogenous increase of size Δ at time T and that no other exogenous changes will occur. Usual arbitrage arguments imply that the path for any price like $\lambda(t)$ must be continuous at time T. The path followed by the equilibrium in the phase plane starts on a trajectory like t_1 such that at time T it arrives at a point exactly Δ units to the left of the trajectory CE from figure 4, which would have been the equilibrium in the absence of any exogenous change in k. As the economy evolves, it moves along t_1 then jumps Δ units to the right to the trajectory CE at time T. Since $e^{-\delta t}\lambda(t)$ can be interpreted as a time 0 market price for knowledge, a foreseen future increase in the aggregate stock of knowledge causes a time 0 increase in the price for knowledge and a consequent increase in the rate of investment in knowledge. Because of the increasing returns, the private response to an aggregate increase in the stock of knowledge will be to reinforce its effects rather than to dampen them. Since the rate of growth of the stock of knowledge is increasing in the level, this kind of disturbance causes the stock of knowledge to be larger at all future dates. Moreover, the magnitude of the difference will grow over time. Thus small current or anticipated future disturbances can potentially have large, permanent, aggregate effects.

As a comparison with the first example shows, this result requires not only that increasing returns be present but also that marginal utility not decrease too rapidly with the level of per capita consumption. If we had restricted attention to the class of bounded, constant elasticity utility functions, $[c^{(1-\theta)} + 1]/(1 - \theta)$ with $\theta > 1$, this phenomenon would not be apparent. The specific example here uses linear utility for convenience, but similar results will hold for constant elasticity utility function $[c^{(1-\theta)} - 1]/(1 - \theta)$ for values of θ close enough to zero.

C. *Example 3*

The analysis of the previous example suggests a simple multicountry model with no tendency toward convergence in the level of per capita output. Suppose each country is modeled as a separate closed economy of the type in example 2. Thus no trade in goods takes place among the different countries, and knowledge in one country has external effects only within that country. Even if all countries started out with the same initial stock of knowledge, small disturbances could create permanent differences in the level of per capita output. Since the rate of growth of the stock of knowledge is increasing over time toward an asymptotic upper bound, a smaller country s will always grow less rapidly than a larger country l. Asymptotically, the rates of growth $(\dot{k}/k)_s$ and $(\dot{k}/k)_l$ will both converge to α, but the ratios k_l/k_s and c_l/c_s will be monotonically increasing over time, and the differences $k_l(t) - k_s(t)$ and $c_l(t) - c_s(t)$ will go to infinity.

It is possible to weaken the sharp separation assumed between countries in this discussion. In particular, neither the absence of trade in consumption goods and knowledge nor the sharp restriction on the extent of the externalities is essential for the divergence noted above. As in the Arrow (1962) learning-by-doing model, suppose that all knowledge is embodied either in physical capital or as human capital. Thus k denotes a composite good composed of both knowledge and some kind of tangible capital. In this embodied form, knowledge can be freely transported between two different countries. Suppose further that the external effect of knowledge embodied in capital in place in one country extends across its border but does so with diminished intensity. For example, suppose that output of a representative firm in country 1 can be described as $f(k, K_1, K_2) = k^{\nu}(K_1^a + K_2^b)$, where k is the firm's stock of the composite good, K_1 and K_2 are the aggregates in the two countries, and the exponent a on the domestic aggregate K_1 is strictly greater than the exponent b on the foreign aggregate K_2. Production in country 2 is defined symmetrically. Then for a specific form of the research technology, Romer (1983) shows that the key restriction on the equilibrium paths Sk_1 and Sk_2 in the two countries comes from the equality of the marginal product of private knowledge imposed by the free mobility of the composite good k:

$$D_1 f(k_1, Sk_1, Sk_2) = D_1 f(k_2, Sk_2, Sk_1). \tag{1}$$

With the functional form given above, it is easy to verify that, in addition to the symmetric solution $k_1 = k_2$, there exists an asymmetric solution. In that solution, if k_1 is larger than k_2 and growing (e.g., country 1 is industrialized and country 2 is not), the path for k_2 that satisfies this equation either can grow at a rate slower than that for country 1 or may shrink, exporting the composite good to the more developed country.[12]

This kind of steady, ongoing "capital flight" or "brain drain" does not require any fundamental difference between the two countries. They have identical technologies. If we assume that there is perfect mobility in the composite k, it can even take place when both countries start from the same initial level of k. If all agents are convinced that country 2 is destined to be the slow-growing country in an asymmetric equilibrium, a discrete amount of the composite good will jump immediately to country 1. Thereafter, the two countries will evolve according to equation (1), with country 2 growing more slowly than country 1 or possibly even shrinking.

This kind of model should not be taken too literally. A more realistic model would need to take account of other factors of production with various degrees of less than perfect mobility. Nonetheless, it does suggest that the presence of increasing returns and of multiple

[12] Details are available in an app. available from the author.

equilibria can introduce a degree of instability that is not present in conventional models. This identifies a second sense in which small disturbances can have large effects. In addition to the multiplier-type effect for a closed economy as described in the last example, a small disturbance or a small change in a policy variable such as a tax rate could conceivably have a decisive effect on which of several possible equilibria is attained.

VII. Conclusion

Recent discussions of growth have tended not to emphasize the role of increasing returns. At least in part, this reflects the absence of an empirically relevant model with increasing returns that exhibits the rigor and simplicity of the model developed by Ramsey, Cass, and Koopmans. Early attempts at such a model were seriously undermined by the loose treatment of specialization as a form of increasing returns with external effects. More recent attempts by Arrow, Levhari, and Sheshinski were limited by their dependence on exogenously specified population growth and by the implausible implication that the rate of growth of per capita income should be a monotonically increasing function of the rate of population growth. Incomplete models that took the rate of technological change as exogenously specified or that made it endogenous in a descriptive fashion could address neither welfare implications nor positive implications like the slowing of growth rates or the convergence of per capita output.

The model developed here goes part way toward filling this theoretical gap. For analytical convenience, it is limited to a case that is the polar opposite of the usual model with endogenous accumulation of physical capital and no accumulation of knowledge. But once the operation of the basic model is clear, it is straightforward to include other state variables. The implications for a model with both increasing marginal productivity of knowledge and decreasing marginal productivity of physical capital can easily be derived using the framework outlined here; however, the geometric analysis using the phase plane is impossible with more than one state variable, and numerical methods for solving dynamic equation systems must be used.[13] Since the model here can be interpreted as the special case of the two-state-variable model in which knowledge and capital are used in fixed

[13] For an example of this kind of numerical analysis in a model with a stock of knowledge and a stock of an exhaustible resource, see Romer and Sasaki (1985). As in the growth model, increasing returns associated with knowledge can reverse conventional presumptions; in particular, exhaustible resource prices can be monotonically decreasing for all time.

proportions, this kind of extension can only increase the range of possible equilibrium outcomes.

References

Arrow, Kenneth J. "The Economic Implications of Learning by Doing." *Rev. Econ. Studies* 29 (June 1962): 155–73.

———. "Applications of Control Theory to Economic Growth." In *Mathematics of the Decision Sciences*, vol. 1, edited by George B. Dantzig and Arthur F. Veinott. Providence, R.I.: American Math. Soc., 1967.

Aubin, J. P., and Clarke, F. H. "Shadow Prices and Duality for a Class of Optimal Control Problems." *SIAM J. Control and Optimization* 17 (September 1979): 567–86.

Baumol, William J. "Productivity Growth, Convergence and Welfare: What the Long Run Data Show." Research Report no. 85-27. New York: New York Univ., C. V. Starr Center, 1985.

Bernstein, J. I., and Nadiri, M. Ishaq. "Research and Development, Spillovers and Adjustment Costs: An Application of Dynamic Duality at the Firm Level." Working paper. New York: New York Univ., 1983.

Blitch, Charles P. "Allyn Young on Increasing Returns." *J. Post Keynesian Econ.* 5 (Spring 1983): 359–72.

Brock, William A. "A Simple Perfect Foresight Monetary Model." *J. Monetary Econ.* 1 (April 1975): 133–50.

———. "A Polluted Golden Age." In *Economics of Natural and Environmental Resources*, edited by Vernon L. Smith. New York: Gordon and Breach, 1977.

Brock, William A., and Gale, David. "Optimal Growth under Factor Augmenting Progress." *J. Econ. Theory* 1 (October 1969): 229–43.

Cass, David. "Optimum Growth in an Aggregative Model of Capital Accumulation." *Rev. Econ. Studies* 32 (July 1965): 233–40.

Cass, David, and Shell, Karl. "Introduction to Hamiltonian Dynamics in Economics." *J. Econ. Theory* 12 (February 1976): 1–10. (*a*)

———. "The Structure and Stability of Competitive Dynamical Systems." *J. Econ. Theory* 12 (February 1976): 31–70. (*b*)

Chipman, John S. "External Economies of Scale and Competitive Equilibrium." *Q.J.E.* 84 (August 1970): 347–85.

Dechert, W. Davis, and Nishimura, Kazuo. "A Complete Characterization of Optimal Growth Paths in an Aggregated Model with a Non-concave Production Function." *J. Econ. Theory* 31 (December 1983): 332–54.

Dixit, Avinash K.; Mirrlees, James A.; and Stern, Nicholas. "Optimum Saving with Economies of Scale." *Rev. Econ. Studies* 42 (July 1975): 303–25.

Ekeland, Ivar, and Scheinkman, José A. "Transversality Conditions for Some Infinite Horizon Discrete Time Maximization Problems." Technical Report no. 411. Stanford, Calif.: Stanford Univ., IMSSS, July 1983.

Ethier, Wilfred J. "National and International Returns to Scale in the Model Theory of International Trade." *A.E.R.* 72 (June 1982): 389–405.

Hansen, Lars Peter; Epple, Dennis; and Roberds, William. "Linear-Quadratic Duopoly Models of Resource Depletion." In *Energy, Foresight and Strategy*, edited by Thomas J. Sargent. Washington: Resources for the Future, 1985.

Helpman, Elhanan. "Increasing Returns, Imperfect Markets, and Trade Theory." In *Handbook of International Economics*, vol. 2, edited by Ronald W. Jones and Peter B. Kenen. New York: North-Holland, 1984.

Hicks, John R. "Thoughts on the Theory of Capital: The Corfu Conference." *Oxford Econ. Papers* 12 (June 1960): 123–32.
Hochman, Oded, and Hochman, Eithan. "Regeneration, Public Goods, and Economic Growth." *Econometrica* 48 (July 1980): 1233–50.
Kaldor, Nicholas. "The Role of Increasing Returns, Technical Progress and Cumulative Causation in the Theory of International Trade and Economic Growth." *Écon. Appl.* 34, no. 4 (1981): 593–617.
Kendall, M. G. *Rank Correlation Methods.* New York: Hafner, 1962.
Kendrick, John W. *The Formation and Stocks of Total Capital.* New York: Columbia Univ. Press (for N.B.E.R.), 1976.
Knight, Frank H. "On Decreasing Cost and Comparative Cost: A Rejoinder." *Q.J.E.* 39 (February 1925): 331–33.
Koopmans, Tjalling C. "On the Concept of Optimal Economic Growth." In *The Econometric Approach to Development Planning.* Amsterdam: North-Holland (for Pontificia Acad. Sci.), 1965.
Kuznets, Simon. "Notes on the Pattern of U.S. Economic Growth." In *The Reinterpretation of American Economic History,* edited by Robert W. Fogel and Stanley L. Engerman. New York: Harper and Row, 1971.
Levhari, David. "Extensions of Arrow's 'Learning by Doing.'" *Rev. Econ. Studies* 33 (April 1966): 117–31. (*a*)
———. "Further Implications of Learning by Doing." *Rev. Econ. Studies* 33 (January 1966): 31–38. (*b*)
McKenzie, Lionel W. "On the Existence of General Equilibrium for a Competitive Market." *Econometrica* 27 (January 1959): 54–71.
Maddison, Angus. "Per Capita Output in the Long Run." *Kyklos* 32, nos. 1, 2 (1979): 412–29.
———. *Phases of Capitalist Development.* New York: Oxford Univ. Press, 1982.
Majumdar, Mukul, and Mitra, Tapan. "Intertemporal Allocation with a Nonconvex Technology: The Aggregative Framework." *J. Econ. Theory* 27 (June 1982): 101–36.
———. "Dynamic Optimization with a Non-convex Technology: The Case of a Linear Objective Function." *Rev. Econ. Studies* 50 (January 1983): 143–51.
Phelps, Edmund S. "Models of Technical Progress and the Golden Rule of Research." *Rev. Econ. Studies* 33 (April 1966): 133–45.
Ramsey, Frank P. "A Mathematical Theory of Saving." *Econ. J.* 38 (December 1928): 543–59.
Reynolds, Lloyd G. "The Spread of Economic Growth to the Third World: 1850–1980." *J. Econ. Literature* 21 (September 1983): 941–80.
Rockafellar, R. Tyrrell. *Convex Analysis.* Princeton, N.J.: Princeton Univ. Press, 1970.
Romer, Paul M. "Dynamic Competitive Equilibria with Externalities, Increasing Returns and Unbounded Growth." Ph.D. dissertation, Univ. Chicago, 1983.
———. "Cake Eating, Chattering and Jumps: Existence Results for Variational Problems." *Econometrica,* vol. 54 (July 1986).
Romer, Paul M., and Sasaki, Hiroo. "Monotonically Decreasing Natural Resource Prices under Perfect Foresight." Working Paper no. 19. Rochester, N.Y.: Univ. Rochester, Center Econ. Res., 1985.
Shell, Karl. "A Model of Inventive Activity and Capital Accumulation." In *Essays on the Theory of Optimal Growth,* edited by Karl Shell. Cambridge, Mass.: MIT Press, 1967. (*a*)
———. "Optimal Programs of Capital Accumulation for an Economy in Which There Is Exogenous Technological Change." In *Essays on the Theory*

of Optimal Growth, edited by Karl Shell. Cambridge, Mass.: MIT Press, 1967. (*b*)

Sheshinski, Eytan. "Optimal Accumulation with Learning by Doing." In *Essays on the Theory of Optimal Growth,* edited by Karl Shell. Cambridge, Mass.: MIT Press, 1967.

Skiba, A. K. "Optimal Growth with a Convex-Concave Production Function." *Econometrica* 46 (May 1978): 527–39.

Streissler, Erich. "Growth Models as Diffusion Processes: II. Empirical Illustrations." *Kyklos* 32, no. 3 (1979): 571–86.

Weitzman, Martin L. "Optimal Growth with Scale Economies in the Creation of Overhead Capital." *Rev. Econ. Studies* 37 (October 1970): 555–70.

von Weizsäcker, Carl C. "Tentative Notes on a Two Sector Model with Induced Technical Progress." *Rev. Econ. Studies* 33 (July 1966): 245–51.

Young, Allyn A. "Increasing Returns and Economic Progress." In *Readings in Welfare Economics,* edited by Kenneth J. Arrow and Tibor Scitovsky. Homewood, Ill.: Irwin (for American Econ. Assoc.), 1969.

Uzawa, Hirofumi. "Optimum Technical Change in an Aggregative Model of Economic Growth." *Internat. Econ. Rev.* 6 (January 1965): 18–31.