# The neoclassical theory of growth and distribution *

ROBERT M. SOLOW

## Introduction

For the sake of clarity, given the history of controversy in this field, it is a good idea to start with some definitions. What is meant by 'neoclassical', by 'growth', and, for that matter, by 'theory'?

'Theory' is perhaps too ambitious a term to describe the subject to be discussed in this article. To call something a theory is to suggest that it aims for an 'ultimate' explanation of fundamental character. The everyday word in economics is 'model'. A model is a simplified representation of a more complex reality. It may be mathematical, but not always. Its claim is to isolate the essential factors, not necessarily to explain them finally. Admittedly the distinction between a theory and a model can vanish at the edges. The immediate antecedent of the neoclassical model of growth is almost always described as the 'Harrod-Domar model' and it might be better to hold to that terminology here. But both words are in common use and that is probably as it should be.

'Growth' means growth of potential output. The idea is to try to isolate relatively smooth, trend-like growth, dominated by supply side factors, from economic fluctuations or business cycles, usually driven by the demand side. There is no implication that either sort of path ever occurs in its pure form in actual economies. (It may be

☐ Massachusetts Institute of Technology, Department of Economics, Cambridge, Mass. (USA).

* This article originated as an entry prepared for the *Storia del XX secolo*, planned and so far unpublished by the Istituto della Enciclopedia Italiana, to whom I express my thanks for kindly authorizing the use of this material.

*BNL Quarterly Review*, no. 215, December 2000.

worth mentioning that the three modern founders of neoclassical growth theory – R.M. Solow, T.W. Swan and J. Tobin – were all 'Keynesian' in their approach to short-run macroeconomics.) This analytical intention takes the form of supposing the available supply of labor always to be fully employed and the existing stock of productive capital goods always to be fully utilized. ('Fully' could be replaced by 'normally' or by any other constant degree of utilization.) This assumption of full utilization could better be made explicit by introducing a government that makes (useless) expenditures and levies (lump-sum) taxes simply in order to preserve full utilization; but this is rarely done, presumably because the financial complications would obscure the essential supply-side orientation of the model. Full employment/utilization is usually just assumed. In other words, saving and investment turn out to be equal at that level of employment and utilization, although the mechanism that brings them into equality is left unspecified.

This is a choice with consequences. It is possible that economic growth and fluctuations are so closely bound together that any attempt to separate them must inevitably omit essential factors governing the growth of potential output. One can imagine a Schumpeterian making such a claim, though its truth is not self-evident. The neoclassical model allows in one important respect for the interaction between fluctuations and growth: fluctuations will surely perturb the rate of investment and that will necessarily affect the path of potential output. There are no doubt other possible interactions, but the neoclassical model ignores them.

Finally, 'neoclassical' is a concept that attracts polemics and, maybe for that very reason, is often unclear. There can be no simple correct definition. Models that are described as neoclassical tend to have some common characteristics, however.

*a*) On the technological side, neoclassical theory is at home with smooth production functions exhibiting constant returns to scale and diminishing returns to individual inputs. Much of the early literature focussed on these technological assumptions, probably in the (false) belief that 'marginalism' requires smoothness. (This literature happened to coincide in time with the domestication of linear programming and activity analysis within economic theory. It has been clear for some time that neoclassical growth theory can operate perfectly well within a technology of discrete activities.) In general, the conclusion is that smoothness and constant returns to scale are assumptions of (great) convenience, but are not essential to neoclassical growth models. Diminishing returns is a much more important matter, as will be discussed later.

There was also much controversy about the role of 'capital' as a factor of production. The ideological overtones of this controversy were pervasive; as a result, it is unlikely ever to be settled to the satisfaction of all participants. Most, perhaps all, of the objections to the use of 'capital' as a factor of production have nothing special to do with the theory of growth; they would apply in any branch of economics. As far as neoclassical growth theory is concerned, 'capital' as an abstract concept or even as a value sum plays no role at all. Potential output does not depend on it. The stocks of capital goods, defined in any necessary detail, do matter, but they are not in any way abstract. There are, however, very many capital goods, especially if they have to be distinguished by age, location and other characteristics. If the neoclassical growth model – or any other growth model – is to be useful in application, there will have to be aggregation over capital goods and over consumption goods and qualities of labor as well. Some part of the controversy no doubt rests on the fact that one side treats this aggregation as a matter of simplification, as with any index number, while the other side regards it as an important matter of principle. Probably the best method of exposition is to think of the neoclassical growth model as being a story about an imaginary economy that has only one produced good that can be consumed directly or stockpiled for use as a capital good. It is then an exact theory of that economy; and it becomes a difficult practical matter whether it provides a useful analogue of a multi-commodity economy. (Later on the obvious extension will come up: to an economy with a single produced capital good and a single physically different produced consumption good.)

One technological requirement remains, if the neoclassical growth model is to be a useful tool: the aggregate economy must be able to function efficiently with a significant range of capital intensities (i.e., overall capital-labor ratios) in production. There is no need for smooth substitutability between capital and labor; it is enough if there are two production activities, each with fixed proportions of capital and labor but one of them significantly more capital-intensive than the other. In practice this may mean only that there are several produced goods with differing capital-intensities; the economy-wide capital-intensity changes with changes in the mix of goods produced. That picture lies outside the one-good model but helps to explain how it can represent a more complex reality. It is doubtful that this can be an important meaning of 'neoclassical' in the growth context.

*b*) The adjective 'neoclassical' can refer to a standard set of background assumptions: that households supply labor and buy goods so as to achieve the highest available level of satisfaction according to a stable set of preferences, constrained by their budgets; that firms make employment and investment decisions so as to achieve the highest feasible profit, in an appropriate long-run sense, constrained by their technological capabilities and the characteristics of factor and product markets; and that all markets tend to clear, mediated by prices.

These assumptions or similar ones are clearly part of the neoclassical tradition, not only in the theory of growth. Some variation is possible. For example, it is often assumed that labor is inelastically supplied in an amount related to the size of the population but independent of the real wage. Then the only remaining household decision is the division of income between consumption spending and saving. The model works perfectly well if the notion of utility maximization is replaced by any reasonable consumption function that relates spending to income, the interest rate and other prices. Similarly, though much less often, the assumption of profit maximization by firms can be replaced by some other systematic criterion of behavior. Market clearing – full employment and full utilization – is a universal – one might say defining – assumption. As mentioned earlier, it is part of the attempt to analyze long-run tendencies independently of short-run fluctuations, to be neoclassical in the long run and Keynesian in the short run.

There has been a characteristic evolution within the neoclassical tradition. The earliest neoclassical growth models treated household behavior by means of a conventional saving function or consumption function. Later on, beginning with papers by Cass (1965) and Koopmans (1965), an alternative convention became established: households were endowed with an intertemporal utility function, almost always a time-additive function of current consumption and perhaps leisure. It is then assumed that current consumption (and perhaps labor-supply) is always the first step in an infinite-time-horizon optimization of utility under perfect foresight about all future incomes and prices. Symmetrically firms are supposed at each step to optimize as price-takers.

Under favorable conditions, the maintained competitive equilibrium in this model is a solution of the 'Ramsey problem'. The model economy traces out just the path that it would follow if it were planned by the single, immortal household, solving its infinite-time utility maximization constrained only by given technological possibilities and the inevitable identities. Under these favorable assumptions, the decentralization to competitive firms does not matter; in effect the industrial sector faithfully carries out the wishes of the household.

Ramsey had formulated exactly this problem as a normative exercise. He thought of himself as working out what an omniscient and omnipotent policy-maker would do when faced with the need to make an infinite sequence of decisions about the allocation of output between current consumption and saving-investment for the future. The utility function was viewed as a social-welfare function. (Ramsey, 1928, argued that the discounting of future utility was inappropriate precisely because it would discriminate unfairly between generations. If the optimizing agent is regarded as a representative household rather than as a trustee, this argument becomes weaker.) The replacement of mechanical consumption and labor-supply functions by explicit intertemporal utility-maximization has some advantages, although it may strain the imagination. In any case, the basic conclusions of neoclassical growth theory are not altered by this amendment. It becomes a different matter when the Ramsey approach is extended to the short run as well, so that every observation is to be interpreted as a point on an optimal path. A neoclassical growth theorist can easily find this extension unacceptable. But that is a different story.

*c*) Still another meaning associated with the neoclassical tradition is the habit of thinking in terms of perfect competition as the standard market form for labor, capital goods and consumption goods. This is surely empirically correct; most neoclassical growth theory does exactly that.

The more interesting question is whether this habit is in any way necessary for the consistent working-out and application of the model. The answer is that it is not. From the very beginning it has been pointed out that imperfectly competitive market forms can be inserted in a fairly straightforward way. Of course a model economy with only one produced good and one primary factor is not exactly a very rich laboratory for thought experiments with alternative forms of industrial organization. More recently, as the theory has been extended by Romer (1990) and others to cover the endogenous production of new technology, one or another form of monopolistic competition has become a necessary part of the model; otherwise there would be no way for innovators to collect the rents that motivate research and cover its costs.

*d*) This last observation is a reminder that another use of the adjective 'neoclassical' has become common. The word is sometimes used to refer to the custom of treating technological change as an exogenous force, something that just happens, regularly or irregularly, without any economically motivated action. With few exceptions, neoclassical growth theorists had nothing to say about the endogenous generation of new technology. No doubt they knew perfectly well that profit-seeking firms incur research-and-development costs and often succeed in inventing new products and improved production processes. It is just as obvious that technological advances can often be understood after the fact as responses to perceived opportunities. That is not the same thing as making the level of technology and its rate of change into routinely determined endogenous variables.

That is what many recent authors have tried to do. This rapidly expanding literature will be described later. Apart from this one difference, most of the newer literature belongs clearly in the neoclassical tradition. A more complicated issue will be discussed in its place. The 'new' growth theory, in some but not all versions, has abandoned the assumption of constant returns to scale in favor of increasing returns to scale, and has abandoned the assumption of diminishing returns to capital goods in favor of constant (or increasing) returns to capital goods. Although there is some misunderstanding about this, the two cases are quite different. The neoclassical model can accommodate increasing returns to scale comfortably. (Perfect competition is then no longer a viable market form, if the increasing returns are 'internal', but it can be replaced by monopolistic competition.) The conclusions of the theory change only in detail. Non-diminishing returns to capital – the factor of production that can be accumulated – is altogether a more important matter. Models of growth with constant or increasing returns to capital behave quite differently and should perhaps not be classified as neoclassical even if they are traditional in other respects. More detail will be added in the proper place.

## Structure of the model

In accordance with what has already been said we are to imagine an economy which knows how to produce a single homogeneous commodity. The inputs into production are limited to labor and (the services of) an accumulated stock of the produced commodity itself (i.e. 'capital'). The supply of labor is given exogenously and may be changing (typically growing) as time passes. The existing stock of capital may depreciate. The usual assumption is that depreciation occurs with the passage of time and not dependent on use; this is not important in the context of full utilization. It is further assumed, almost universally, that the amount of depreciation is proportional to the existing stock; the importance of this simplification is that the amount of depreciation does not depend on the age-composition of the stock.

Under the rules of the game, all (or a fixed fraction) of the available labor and existing capital are used in production, and this determines, through the existing technology, a level of output or income, gross of depreciation. One can imagine total income as being distributed to participants in the economy according to whatever institutional rules prevail. The usual (neoclassical) story is that a competitive labor market will set the wage in terms of product equal to the marginal product of labor at current input levels. If production is characterized by constant returns to scale, the most favorable environment for competition, the remaining income will just cover the marginal product of the existing stock of capital. But at this basic level, there is no need to be so specific. Any systematic rules for describing the distribution of gross or net income will do.

At this point the critical economic (or allocation) decision must be made. Gross or net output has to be divided into a part that will be consumed currently and the remainder that will be added to the stock of capital goods. The outcome may depend on the distribution of income or it may not. One may imagine more or less mechanical rules for the determination of consumption or investment, but the rules of the game require that these always add up to total output. Alternatively one may imagine more or less mechanical optimization procedures to determine the current allocation between consumption and saving/investment. However that is arranged, the model economy is now ready to enter the next period or instant. The new supply of labor is given; the new stock of capital is known, after accounting for the previous period's depreciation. The whole process can begin again. In this way the model economy traces out, step by step, its growth path. A change in the allocation rule, or in the distribution rule if allocation is sensitive to distribution, will alter the growth path in a determinate way.

It may be worth mentioning explicitly that it is possible to allow the supply of labor at any instant to depend on the real wage. At any instant, given the inherited stock of capital goods, the demand for labor can be taken as a function of the product wage, usually a decreasing function. Obviously the setting of perfect competition in the product market is especially congenial to the neoclassical tradition. But other institutional assumptions are possible, although they may involve complications with expectations and strategic behavior. All that matters is that there be a determinate relation connecting the real wage and employment. It and the supply curve of labor will together pin down the real wage and the level of employment. The presumption of full employment is replaced by the presumption that workers are 'on their supply curves'.

## Steady states

The model just described would lend itself to iterative computation of the evolution of the economy and therefore to calculation of the effects on that evolution of alternative initial conditions, alternative parameters and alternative rules of behavior. It is more usual, however, to make further simplifying assumptions instead which then allow simpler comparative statements. In particular, suppose production exhibits constant returns to scale, the labor force grows exponentially (so that it is characterized, at least for long intervals, by a constant rate of growth), and technology is stationary.

Then the model is capable of tracing out a 'steady state'. That means a path along which the supply of labor, the level of output and the stock of capital all grow at the same rate. It then follows that savings and investment are always a constant fraction of output. Steady states are simple to describe and compare because they can be characterized by a few parameters. The common rate of growth is just the exogenously given growth rate of the labor force. The other interesting parameters have to be determined from the model. They include the ratio of the stock of capital to employment (a measure of capital intensity), the ratio of output to employment (a measure of productivity) and the ratio of saving and investment to output. They are connected by a necessary relation: the product of the saving rate and productivity must equal the product of the growth rate and capital intensity.

Such a steady state has an obvious deficiency as a description of a modern capitalist economy. A model that gives rise to such paths must be equally defective. The history of modern industrial economies suggests that productivity usually rises, and so does capital intensity. In fact they rise at approximately the same long-run rate, because there is little or no trend in the ratio of capital stock to output. This defect can be eliminated by dropping the implausible assumption that technology is stationary.

Suppose instead that technological progress occurs. At this level the theory has nothing to contribute to an explanation of technological progress, so it is taken as an exogenous fact. Moreover the advance of technology affects production in a very special way: as if the intrinsic effectiveness of each unit of labor grows exponentially at a rate that is approximately constant for substantial intervals of time. This is described as 'labor-augmenting' or 'neutral' (or more specifically Harrod-neutral) technological progress. If the labor force grows at 2 percent per year and the rate of labor-augmenting technological progress is 1 percent per year, then the 'effective' labor force grows at $2+1=3$ percent per year. This specialized notion is needed only to allow for a slightly generalized sort of steady state.

In the definition of a steady state given earlier, it is only necessary to replace 'labor force' or 'employment' by 'effective labor force' or 'effective employment' and leave everything else the same. Then, in this new sort of steady state, output and the stock of capital grow at the same rate as the effective labor force, i.e., at a rate equal to the sum of the population growth rate and the rate of technological progress. The ratio of output to effective employment is constant in a steady state, so productivity – output per worker in natural units – is growing at the same rate that technology improves. This is of course the quantity that matters for economic welfare. Capital intensity behaves similarly, so the capital-output ratio is still trendless, and so is the ratio of saving and investment to output.

This sort of steady state plays a large role in neoclassical growth theory, but it is not a neoclassical invention. Nicholas Kaldor's account of the "stylized facts" of capitalist growth (1961) describes exactly such a steady state. (Kaldor added the trendlessness of the share of wages in the national income; neoclassical steady states will have this characteristic too if wages are determined competitively, or if the degree of monopsony in the labor market and the degree of monopoly in the product market are merely non-zero but constant.)

In summary, the neoclassical growth model can get along without (generalized) steady states. The advantage of the concept is that it appears to be empirically relevant and analytically simple. The habit of talking in terms of rates of growth is certainly convenient. By now it is hard to imagine any comparably convenient way to describe long time series. No doubt one could learn to think in terms of more complicated time paths than exponentials, if they were empirically useful; and eventually one would work out what assumptions about the building-blocks of the theory give rise to just those paths.

The main disadvantage of the habit follows from what has just been said. An uncritical focus on steady states leads easily to the unthinking adoption of just those assumptions about building-blocks that give rise to them, even if the assumptions are otherwise implausible. Each vice reinforces the other. One way to retain the simplicity at lower cost might be to insist that steady-state-like paths are episodes; they do not last forever. A growth path may consist of segments, each rather like an exponential; when such segments are spliced together, the result need not be exponential at all.

In the contemporary world of cheap computation for the simulation of dynamic processes, the sorts of questions one wants to ask, about the effects of changing basic assumptions on the nature of paths, can be asked and answered without the requirement that the paths be restricted to some simple parametric family. There is no need to look for steady states unless the empirical material suggests it is a good idea. There is certainly no need to reject otherwise convincing assumptions because they make steady-state analysis difficult or impossible.

## The basic question

Perhaps the best way to understand the working of the neoclassical growth model is to ask a simple 'practical' question. Imagine a fully aggregated one-commodity economy of the kind already described, and suppose that it is and has been for some time in a steady state of the kind already described. Now further imagine that at some instant of time there is a once-for-all increase in the fraction of output invested, say from 10 to 12 percent, and then the rules of the game take over. What will the new growth path of the economy look like? How will it differ from the steady state path that would have continued if the rate of saving-investment had not changed?

In this version of the model, evidently the saving-investment rate is taken as a parameter. If instead the model were one in which a representative household were maximizing an additive intertemporal utility, one would have to imagine a sudden maintained change in one of the parameters of the utility function, like the rate of time preference or the elasticity of substitution between current and future consumption. The qualitative answer to the question provided by the model will be the same, although the details may differ.

In the initial period, aggregate output is unchanged. The stock of capital is what it was, and the assumption of full employment says that employment will be whatever it would have been without the increase in investment. All that changes is the composition of output: more is invested and less is consumed. In the next period, therefore, aggregate output will be higher than it would have been if the original steady state had been maintained, because the capital stock will be higher and employment the same. So now investment is larger for two reasons: output is larger and the fraction invested is higher. Evidently the model tells us that the new path for aggregate output and capital stock will lie above the continuation of the old steady-state path. Since the two paths started in the same place, the new path must have a higher growth rate.

The characteristic conclusion of the neoclassical growth model is that the new path for output and capital stock will continue to lie above the old one, but the higher growth rate will be temporary. The new path gains on the old one only because it has higher and increasing capital intensity. There is more capital but the same amount of labor. Diminishing returns to capital goods implies that each increment to capital intensity brings a smaller increment to output. Under the standard assumptions of the theory, the new path approaches a steady state. This steady state has higher capital intensity and therefore higher productivity than the old steady state. But its growth rate is the same: it is equal to the sum of the rate of growth of employment and the rate of labor-augmenting technological progress, and that has not changed.

There is an analogous exercise: starting from a steady state, imagine a sudden maintained increase in the growth rate of labor force and employment. In the initial steady state, investment is just large enough to equip each period's increment to employment with capital at the constant economy-wide degree of capital intensity. If the saving-investment rate is unchanged, this will no longer be so when employment is growing more rapidly. Capital intensity – and productivity – will gradually fall until the old rate of investment is capable of maintaining it even with more rapidly growing employment. Thus a steady state is re-established with rate of growth equal to the new, higher population growth rate, and lower capital intensity and productivity. (If there is labor-augmenting technological progress, a slightly more complicated statement is true.) In the new steady state the aggregate growth rate of capital and output will be higher, but rates of growth of capital and output per worker in natural units will be equal to the unchanged rate of labor-augmenting technological progress.

## Convergence to steady states

The stories just told are comparisons between steady states. But they have been told in a way that presumes something stronger: that the model economy, if it starts outside of a steady state, will move toward one and eventually arrive there. In the simplest typical versions of the neoclassical growth model, that presumption is true. There is a unique steady state and the model converges to it from any initial position. The stories just told give a rough indication of the economic reasons for this conclusion.

In more complicated versions of the model, this neat conclusion need not hold. The complications can arise in several ways. For example, if the saving-investment rate is made a function of economic variables – return on capital, distribution of income – there may be more than one possible steady state. The same thing can happen if the population growth rate varies with, say, the level of consumption per head. It is easy to find plausible assumptions that allow for a low-level (Malthusian) steady state and a high-level steady state (after a 'demographic revolution'). As is usual in such situations, where there are two stable steady states they are separated by an unstable steady state. In these cases it remains true that the model economy always tends to a steady state, no matter where it starts; but now the particular steady state to which it converges depends on the initial position of the economy.

Another possible source of instability, and perhaps of a multiplicity of steady states, comes from non-convexities in production, i.e., from regions of increasing returns to capital. More will be said about this later; Azariadis and Drazen (1990) have a model in which, as capital intensity increases, there are occasional discrete jumps in productivity as the economy enters a new 'regime' characterized by a different organization of production.

When that happens, there may be many steady states, alternatively stable and unstable. If there are diminishing returns to capital within any regime, then very likely the model economy will converge to some steady state no matter where it starts.

Versions of the neoclassical model that incorporate infinite-horizon intertemporal optimization by the representative consumer are subject to a quite different possibility of non-convergence. Generally such models have a unique steady state and generally the optimizing trajectory converges to it. The conditions that define an optimal trajectory fall into two groups. There are local conditions that tie together neighboring instants of time; and there is a 'transversality condition' that requires and exploits long-distance perfect foresight. There are many trajectories that satisfy the local conditions. One of them satisfies the transversality condition and it is a stable trajectory, converging to a steady state. The other paths, satisfying only the local conditions of optimality, do not converge to any steady state; in fact they may bring the model economy to some nonsensical state of affairs, by decumulating capital, for instance. Local optimality is easier to achieve than anything that depends on foresight about the distant future. One would have to ask how an economy would 'know' that it is on a locally optimal but globally unstable path. This sort of ('saddle-point') instability is a negative feature of models involving intertemporal optimization (on this, see Dixit 1990).

Still another sort of non-convergence arises when there are many kinds of capital goods and investment has to be allocated among them. This was first uncovered by Hahn (1966); it is strictly speaking outside the one-commodity model discussed here, but it reveals a potential difficulty that should not be ignored. This difficulty is extensively discussed in Burmeister (1980). Informed speculators might push the economy back toward the convergent path from time to time; but this stabilizing force would require foresight-at-a-distance, as discussed above.

## Speed of convergence

It was very early recognized that the convergence of the standard neoclassical model to a unique steady state (supposing there to be one) would normally be very slow. (Atkinson 1969 is a typical example; there are others.) Imagine the model economy in its steady state and suppose that half the capital stock is suddenly destroyed by war or earthquake. Calculations with empirically reasonable shapes of functions and values of parameters show that restoration of the initial position to within a percent or two would occur only after the lapse of a century or more. In effect the normal process of economic development has to be retraced at the same speed as the original trajectory.

More recently King and Rebelo (1993) have looked at the same problem in the context of a model with long-time-horizon intertemporal maximization. The picture is then quite different. Restoration of the initial position occurs quite quickly. It is easy to see why: the foresighted optimizing consumer responds to the loss of initial capital by a temporarily higher rate of saving and investment. The earlier stages of restoration occur quickly and saving-investment is gradually reduced as the steady state is approached.

The point was well understood by early theorists, although it does not seem to have been stated in the published literature. The reasoning is straightforward even without formal optimization. Loss of a substantial fraction of the stock of capital goods would, with full employment, lead to a substantially higher return on capital and very likely to increased investment. In simple numerical experiments, the saving-investment rate is assumed to jump by some fixed amount and remain there until the return on capital has risen to a level closer to its original steady-state value, when investment reverts to normal. The result is, of course, qualitatively similar to King-Rebelo's (1993). The restoration time is drastically reduced.

## Primary factors

Much of the characteristic flavor of the neoclassical model derives from the role of labor as a primary factor of production, i.e., an essential factor of production that is not producible within the model. For instance, the rate of growth of the primary factor is the steady-state rate of growth for the economy as a whole in the absence of technological progress. Even with technological progress, the rate of growth is given by the growth rate of the 'effective' primary factor. It will be seen later how the model is changed drastically when there is no primary factor at all. Here it is worth mentioning the possibility of more than one primary factor.

There are two obvious candidates for the role of second primary factor. The first is a second quality of labor, supplied according to its own autonomous rules and only imperfectly substitutable for the first quality. The other is a natural resource, either a self-renewing resource like land, with growth rate zero, or an initial stock of an exhaustible resource like oil, that is used up over a very long time. (This is an inadequate description of the supply of most non-renewable resources, but this is not the place to go into detail.) Only a few remarks are in order here.

What can be said about the steady-state growth rate for a model economy with two such primary factors, in the absence of technological progress? If land were the only primary factor, the growth rate would have to be zero, because at this level of abstraction land functions like a constant labor force. The same would be true if the single primary factor were 'oil' with the additional complication that the steady-state level of output might have to be zero as well. (That is one possibility. It is also possible, if 'oil' and capital are good substitutes, that a zero-growth steady state with positive consumption can be maintained.)

Evidently, when there are two or more primary factors, a lot depends on the nature of the technology for the production of final output. To take an extreme case, if production can be carried on without land or 'oil' – using only capital and labor – then the existence of a second primary factor makes no essential long-run difference. One can imagine less extreme cases in which the second primary factor is indispensable for final production in the sense that zero 'oil' means zero output, but a given level of output can be produced with smaller and smaller inputs of 'oil' but larger and larger inputs of capital. Cases can be exhibited (Stiglitz 1974) in which the steady-state rate of growth can be anything between zero and the population growth rate, depending on the rate of capital investment. This is one way in which the neoclassical model allows the growth rate to depend on the investment rate; but it is obviously very special.

Whether there is one essential primary factor or more than one, the lesson is the same in this class of models. 'Permanent' growth can escape the limitation imposed by primary factors only by virtue of technological progress. At the most basic level, the reason for this conclusion is the assumed presence of diminishing returns to producible factors of production. This result is not very informative in practice. To mention only two reasons, it is not always clear whether any particular input is 'essential' or indispensable. Energy is no doubt essential, but energy from fossil fuels is surely not; and furthermore 'permanent' is just a dramatic word for 'long-lasting' in this context. One imagines that conclusions valid for infinite time are approximately valid for very long but finite times, although this is rarely made precise.

## 'New' growth theory

The neoclassical model, as described so far, leaves an obvious gap. The long-term or steady-state rate of growth itself is governed wholly by forces that the model treats as exogenous. Those forces are the growth rates of primary factors and the rate of technological progress. To treat something as exogenous is not the same thing as to believe that it actually is determined by non-economic or non-social forces. The implication is only that those forces are not determined within the model, either because their determination is not well understood or because the focus is on the interaction of those forces with others that are endogenous to the model.

For example, it is customary for neoclassical growth models to take the rate of growth of population and labor force (and therefore employment) as given, although everyone would agree that demographic trends are responsive to economic growth itself. From the very beginning (Solow 1956) examples have been given of the way demographic and economic growth could interact within the theory, and an occasional article (Johansen 1967) has been devoted entirely to this theme. But as long as the current ideas about endogenous population growth did not go significantly beyond the Malthusian scheme, little or nothing was lost by assuming population growth to be exogenous. The model is then answering questions about the long-run economic effects of demographic changes. The reverse implications could be filled in informally or formally if that were useful.

Much the same thing could be said about endogenous technological progress. It is an old idea with a literature of its own (for early examples, see Fellner 1961, Kennedy 1964, Samuelson 1965 and

Weizsäcker 1966). But for a long time neoclassical growth theorists seemed to have nothing to add. Once again there were important exceptions, for example Arrow's (1962) concept of "learning by doing". It is interesting that Arrow's use of this idea implied no great change in the general structure of the neoclassical growth model, with the important qualification that it provided a well worked out mechanism that creates a systematic divergence between perfectly-competitive growth and socially optimal growth. (If, for instance, gross investment is the vehicle for productive learning, private optimization will lead to underinvestment. Private investors will not take account of the external benefit they induce.)

Beginning in the middle 1980s with the work of Romer (1986) and Lucas (1988) and continuing with a flood of articles there has been an active attempt to extend the neoclassical model by making the steady-state rate of growth itself an endogenous variable. In its first phase, 'endogenous growth theory' worked by trying to find plausible assumptions that would deny the existence of diminishing returns to the class of productive inputs that can be accumulated by some form of saving and investment. Another way to describe this line of thought is to say that it tried to create a model with constant returns to scale but no primary factors at all. It is fair to say that this first phase failed, for several reasons. The second phase of endogenous growth theory worked by trying to construct interesting and plausible models of the generation of technological progress as a normal profit-seeking economic activity. That is to say, it has tried to endogenize technological progress itself. This effort will stand or fall by virtue of the empirical plausibility and utility of its view of research and development and innovation. The books are not yet closed on that issue.

*First phase.* In the one-good context, suppose that output is simply proportional to the stock of same-good capital. Labor is not required for production. Now add any mechanism that will make net investment proportional to output or income, all the time or merely in a steady state. Obviously, then, net investment is proportional to the stock of capital; the factor of proportionality is just the product of the ratio of output to capital and the ratio of saving-investment to output.

This combined proportionality factor is clearly the rate of growth of capital and, therefore, of output. The steady-state (at least) growth rate is determined as the product of the productivity of capital and the saving-investment quota. (In its most straightforward form, this relation was studied by Domar 1946. Harrod 1939 preceded him, but in a more complicated context, mixing short-run and long-run considerations.) What has happened here is that the assumption of diminishing returns to capital has been replaced by constant returns to capital. The drastic consequence is that now the steady-state growth rate depends proportionally on the fraction of income saved and invested, whereas before the growth rate was independent of the investment quota. The growth rate is endogenous. Anything – the amount and nature of taxation, the allocation of public spending, etc. – that can permanently affect the rate of investment can permanently affect the rate of growth. This is not a business-cycle result; it holds with a constant rate of capacity utilization.

It goes without saying that the absence of labor as a factor of production is utterly unrealistic. But this defect can be softened. Arthur Lewis (1954), thinking about economies in the earliest stage of industrial development, assumed that there was an unlimited supply of labor, meaning that employment could be increased freely at a constant real wage. As long as that situation lasts, constant returns to scale is equivalent to constant returns to capital; whenever the stock of capital is increased by x percent, employment is increased by x percent through the absorption of surplus labor, and thus output increases by x percent as a result of constant returns to scale. Everything thus proceeds exactly as if there were constant returns to capital alone, until the reserve of surplus labor disappears. A non-scarce primary factor is like no primary factor at all. Few development economists today are willing to accept this picture of the process.

Another way to make this sort of model more plausible is to extend it to a many-good context. Suppose there is a group of capital goods that can be produced under constant returns to scale, using as inputs only themselves. There are thus constant returns to the group of inputs that can be accumulated, and this is the essential feature. To this sub-economy can be added a consumer good (or more than one) that is produced in the 'normal' way by labor and the bundle of capital goods under conditions of constant returns to scale and diminishing returns to capital and labor. In a steady state, the growth rate of consumption is an average of the growth rates of labor and the composite capital. Since the growth rate of capital is sensitive to the saving-investment quota, as in the Domar model, so is the growth rate of consumption. That solves the problem formally; but the existence of a capital-goods sector using no labor or other primary factor remains a fundamental implausibility.

The most plausible way to tell a story involving constant returns to capital is to think of the second capital good as 'human capital' or even 'technological knowledge' that can be produced via some simple technology of its own, using raw labor and human capital as inputs. That construction seems over-simple, but clearly acceptable. The difficulty arises because somewhere, somehow, an assumption of constant returns to capital alone must be inserted. To take just one example, in Lucas's very influential paper (1988), it is assumed that each period's accumulation of human capital is proportional to the product of the existing stock of human capital and the amount of raw labor devoted to training. There are thus constant returns to human capital in the production of human capital. The rest follows easily. The rate of growth of the stock of human capital is proportional to the amount of raw labor undergoing training. Anything that affects that quantity – and it is a *level*, not itself a rate of growth – will affect the growth rate of human capital and thus the growth rate of output as a whole.

There are other ways of achieving this outcome but they all share this special characteristic. It is worth pointing out explicitly that the special characteristic is constant returns to capital, not increasing returns to scale. It is possible to have increasing returns to scale together with diminishing returns to each factor of production separately. In such cases, the growth rate is not endogenous. The question of increasing returns to scale arises only because the combination of constant returns to capital and the existence of any non-capital factor of production with positive marginal productivity implies increasing returns to scale. But that is merely incidental.

This way of endogenizing the rate of growth – generically known as 'AK models' to emphasize the proportionality of output to capital – has not been a success. There are at least two sets of reasons for this, one theoretical and the other empirical.

The theoretical deficiency is that these models are the exact opposite of robust. They require exactly constant returns to capital. If there is the slightest touch of diminishing returns, then the model becomes standard-neoclassical and does not deliver an endogenously determined rate of growth. There is no slack on the other side: if there is the slightest touch of increasing returns to capital, then the model becomes too powerful for its own good and generates infinite output in finite time. This problem cannot be escaped through the presumption that a very small amount of increasing returns would postpone the onset of infinite output to so distant a future as not to matter. Under the usual assumptions, the appropriately small degree of increasing returns would have to be so small as to be in any practical sense indistinguishable from constant returns. The AK model can survive only if there are exactly constant returns to capital and there is no reason in principle why any such thing should be true.

The empirical evidence for constant returns to capital is, at best, ambiguous. Most of the relevant time series have rising trends in industrial economies. It is then inevitably difficult to distinguish the effects of increasing returns from those of purely time-dependent technological trends. It would seem, at this stage, that the weight of the evidence goes against the AK model. The emphasis on human capital is confirmed by empirical studies. But the model's insistence on constant returns to the complex of capital goods, including human capital, does not seem to be consistent with the historical data for advanced industrial economies. A paper by Mankiw, Romer and Weil (1992), for example, leads the authors to the conclusion that the elasticities of real GDP with respect to tangible capital goods, human capital and raw labor are each about one-third. Thus the combined elasticity for tangible and human capital – centered around two-thirds – is noticeably less than the figure of unity needed by this version of endogenous growth theory. It is difficult to define, and even more difficult to measure, the input of human capital. Other attempts have cast some doubt on the Mankiw-Romer-Weil estimates, but their qualitative conclusion seems to stand.

*Second phase*. In a series of papers that is still continuing, Romer (1990) and many other authors have extended the neoclassical growth model to include the direct, endogenous generation of the rate of technological progress. In such models there is an activity that transforms labor, capital and other resources into improved technology. There is, in effect, a known technology for creating new technology. This activity is operated either by goods-producing firms or by a special class of firms that sell new technology to goods-producing firms.

The long-run rate of growth is genuinely endogenized as part of the normal operation of profit-seeking business firms.

This way of proceeding brings a number of intellectual advantages. It permits the modelling of a basic idea of Schumpeter's: that access to a new technology confers at least temporary monopoly power on the first user. In some models (Aghion and Howitt 1992, for example) the appearance of a new technology makes earlier technologies obsolete; this is Schumpeter's "creative destruction" made precise.

It is inevitable in this line of thought that some form of monopolistic competition must be the normal market form. If production is carried out under constant returns to scale with respect to the standard inputs – labor, services of capital goods, intermediate materials, etc. – then perfectly competitive imputation will exhaust the product and leave nothing over to pay for the resources used in the generation of innovations. The appearance of monopolistic competition is thus necessary. It is also a desirable step toward realism. Research along these lines is continuing actively.

In the end, the success of these models of endogenous technological progress must be judged in terms of the plausibility and empirical success of their representation of the process of innovation. Endogenous growth theory has so far made little contact with those economists who are engaged in the micro-study of innovation (Rosenberg 1982, Nelson 1981). There is an inevitable discordance here. And it is not unique to growth theory; it occurs everywhere in economics. Those who study the innovation-process in great detail, in historical examples or in the study of decision processes in industrial laboratories, will focus on subtle details. However true and interesting these are, they cannot be reduced to the mechanical simplicity that is required if they are to be incorporated into a model of economic growth. Model-builders, on the other hand, will experiment with simple and clever formulations that permit them to get on with model-building. The first group will say that the second group lacks any refined understanding of the process. The second group will say that the first group lacks any refined understanding of what theory is all about. There is no resolution of this conflict. One hopes that it leads to a spiral and not a circle.

Even from the model-building point of view, the endogenization of technological progress is not at all straightforward. Important analytical choices are made for convenience, and then as a matter of habit. In the nature of the case it is very difficult to evaluate their empirical relevance; there are few tested generalizations about the creation of new technology.

Essentially all second-phase endogenous growth models make use of the following idea. Suppose that the current 'level of technology' can be represented by a number *T*. A higher value of *T* corresponds to a more productive technology. In the sector of the economy that produces new technology, assume that the allocation of a certain level of resources (*R*) in each period will result in an increase in *T* in each period. This is no doubt over-mechanical, but any theory of innovation will have to say something very similar. Now comes the crucial decision. If the use of resources *R* for one period transforms *T* into $(1 + g)T$ in the next period, where *g* is of course a function of R, then we have an endogenous theory of growth. Continued allocation of *R* to innovation will generate growth at rate $g(R)$. Anything that motivates a sustained increase or decrease in *R* will result in faster or slower steady-state growth. It is very easy to imagine policies that change *R*: subsidies to research and development, improvements in technical education, etc.

On the other hand, assume instead that a one-period allocation of resources *R* will transform *T* into $T + h$, where now *h* is a function of *R*. Then steady allocation of *R* will not lead to steady growth. It would take a continuing increase in *R* to change the rate of growth, and a permanent increase in the rate of improvement of technology might be impossible to achieve in this way. In this case there is a theory of endogenous technological change, but it does not lead to a theory of endogenous growth. The difference rests only on whether a given level of effort in research is capable of creating a proportional change in the level of technology or an absolute change in the level of technology. There is in fact very little basis for choice. In that sense, endogenous growth theory is still some distance from a plausible and usable model.

This is still an active field of research. Many economists, notably Paul Romer (1986), Alwyn Young (1993) and Gene Grossman and Elhanan Helpman (1991), are developing models of endogenous technological progress that try to capture more of the substance of particular innovations: quality ladders, learning phenomena, clustering and imitation, etc. These alternative ways of thinking about innovation have only slightly different implications for growth. They must all in any case be reduced to a simple form in order to be incorporated in a model. Progress may come from this work in two ways. The greater contextual detail offers some promise of connecting up with micro-studies of the research and development process. Even beyond that, it may lead to empirical work that will be able to throw light on the important question whether a sustained change in the level of innovative activity can generate a sustained increase in the steady-state rate of growth.

## Extensions of the basic framework

### *The two-sector model*

The obvious way to go beyond the completely aggregative character of the standard neoclassical model is to allow for distinct sectors producing consumer goods and investment goods, each with its own technology. The existing stocks of capital goods and labor are then allocated between the two sectors and simultaneously a full set of prices is determined to equilibrate markets for the two goods. (Capital goods may be non-shiftable once purchased, but this will make little difference in or near steady states.) Current production of investment goods is added to the capital stock, and depreciation deducted. This small general-equilibrium model can then be solved period by period. For this extension to make a substantive difference as compared with the one-sector model, there must be a systematic bias in factor-intensity between the consumer-goods sector and the sector producing investment goods.

The literature begins with a paper by Uzawa in 1961 and soon occupied a substantial group of theorists. A good textbook treatment is to be found in Burmeister and Dobell (1970) and there is a definitive treatise by Duncan Foley and Miguel Sidrauski (1971). The steady-state properties of the two-sector model turn out not to be very different from those of the one-sector model. In other respects, however, some interesting insights emerge.

The relative factor-intensities of the two sectors play a role in the convergence properties of the model. Generally speaking, stability is favored when the consumption-goods sector is more capital-intensive than the investment-goods sector (in the sense that it uses a higher capital-labor ratio at any common input prices, i.e. wage rate and rental price of capital). More generally one can say that convergence to a steady state is much more problematical in the two-sector case than in the one-sector model, where it is hardly problematical at all.

Another set of interesting questions opened up by this extension has to do with the 'incidence' of technological change. In the one-sector model, the only issue is whether technological change augments (or 'saves') labor or capital or both in different proportions. In the two-sector case there is a fourfold interaction: innovation can differentially augment labor or capital in the consumer-goods sector, the investment-goods sector, or both. The implications for growth can be quite diverse (and there are applications to trade between growing economies).

The two-sector model provides a richer environment for the study of fiscal policy than the standard one-sector model. This is because tax policies and public expenditure programs (including income transfers) will ordinarily have differential effects on the consumption-goods sector and the investment-goods sector. Foley and Sidrauski (1971) make an extensive analysis of the implications of alternative fiscal-policy actions. So far as the steady-state growth rate is concerned, however, the implications remain neo-classical.

### *Multi-sector models*

One might think that the two-sector model would illustrate most of the complications that could arise in models with many sectors. But this is not really true: multi-sector models introduce some wholly new principles. This is because the intersectoral allocation of resources in the two-sector model still rests on the consumption-saving decision, as in the one-sector case. A true multi-sector model, on the other hand, needs a way of allocating consumption expenditure among different consumption goods, and a way of allocating investment expenditure among different capital goods. If these choices are constrained to follow simple fixed-proportions rules, then the model is not truly multi-sectoral. The only alternative course is to introduce an economically interesting principle of allocation, and this can change the nature of the model substantially.

In the early literature, Gale (1967) studied the pure theory of $n$-sector models, but usually in the special case where the technology of the economy has the fixed-proportions input-output form. Johansen's work (1974) was more practically oriented. He was interested in the formulation of a multi-sector model that could be implemented empirically and used for planning or forecasting. Another early example is Mahalanobis (1953), now thought to be too rigid and special. Luigi Pasinetti (1981) has studied multi-sector models with the goal of uncovering structural patterns connected with growth. These can arise either from the technological side, when input proportions are systematically related to the scale of an industry's output, or from the demand side, when consumption patterns are systematically related to the level of income.

*Overlapping generations*

In the version of the neoclassical growth model that rests on plausible, ultimately empirically based, saving and expenditure functions, there is no necessary demographic structure. Each instant's allocations just occur according to the rules. In the version that rests on an intertemporally optimizing representative-agent, there is only the simplest possible demographic structure: an immortal individual or dynasty with perfect foresight to infinity. If the first version may be thought to lack proper micro-foundations, the second version surely lacks any sort of plausibility.

The overlapping-generations structure, introduced by Samuelson (1958), falls somewhere between the other two. (The literature takes off from Samuelson, but Allais had earlier in 1947 proposed the same model.) The simplest case, and the one almost always used, involves a population that is growing at a constant geometric rate, which may be zero. Each individual lives for two periods; if the population is growing at $g$ percent per period and $N$ individuals are born in a certain period, then $(1 + g)N$ are born in the next period. Thus two generations coexist in each period, the one born that period and the one born the period before.

The pattern of economic activity varies from model to model. Typically a generation supplies labor in its first period of life, earns wages, spends part and saves the rest, investing the savings in whatever vehicles are available. In its second period of life, a generation spends on consumption its earlier savings plus whatever return those savings have earned in the intervening period. A generation may or may not be endowed with labor in its second period. If not, its prior savings (plus interest) are all that it has to spend on consumption when it is old. Most often, successive generations are born without any inherited endowment and leave no bequest; but this can be modified so that the young have an inheritance to spend when they are young, in addition to wages earned, and spend less than all their resources when they are old, leaving part as a bequest. The important point is that each generation has perfect foresight only for two periods (i.e. one lifetime, which is a lot, but less than infinity) and chooses an optimum work-saving-investment-consumption plan only for its own lifetime. This ingenious model has been much used.

The overlapping-generations model connected with mainstream growth theory in a paper by Peter Diamond (1965). Excellent textbook treatments at an advanced level can be found, briefly, in chapter 3 of the book by Blanchard and Fischer (1989) and, much more exhaustively, in books by Azariadis (1993) and Farmer (1993).

In Diamond's formulation (1965), now fairly standard, the young household formulates a two-period life-cycle consumption plan in light of its first-period earnings and the prospective rate of return on its saving. There is then a well-defined saving function: the effect of higher earnings is to increase saving, but a higher rate of return has ambiguous effects, for well-known reasons. This saving function, combined with the usual growth-theory convention that the markets for labor and goods (and securities, if that is how investment is financed) all clear, defines a determinate growth path. It has properties much like those of the standard neo-classical model. There will usually be a non-trivial steady state; under the usual assumptions, there will be only one steady state.

This model has been used by Diamond (1965), Samuelson (1975) and others to study the effects of fiscal policies like debt-financed public consumption and the institution of a social security system. A representative result is that a fully funded social security system has no effects on the growth path, but a pay-as-you-go system reduces national saving and shifts the destination of the path to a steady state with a lower stock of capital.

When the overlapping-generations model is extended to include 'outside' money, with perfect foresight or rational expectations, its character changes rather drastically. In particular, there can be a continuum of equilibrium growth paths starting from the same initial conditions. Even more dramatically, there can be irregular equilibrium paths induced by more or less arbitrary, but self-confirming, expectations. These go under the name of speculative bubbles or 'sunspots'. (The word refers to the possibility that the mere belief that economic activity is correlated with sunspots could cause agents to take actions that would bring about fluctuations that are indeed correlated with sunspots.) These matters, of more concern to the theory of fluctuations than to the theory of growth, were first discussed by Cass and Shell (1983); Farmer (1993) gives an excellent exposition.

*Money and growth*

Up to this point all of the theory that has been discussed has been entirely in real terms. No nominal asset and no nominal price has played a role. (That is not quite true for the overlapping-generations model. Notice the title of Samuelson's original paper. But most of the growth-theoretic part of the overlapping-generations literature is in real terms.) There is, however, an extensive body of work exploring the implications of 'outside' money for the neoclassical theory of growth. This literature has been summarized and evaluated by Orphanides and Solow (1990), which can be consulted for references.

The framework for this discussion and the statement of the main problem is due to Tobin (1955). In static monetary theory, the question of the neutrality of money has been exhaustively discussed and the issues are well understood. That is, as a matter of comparative statics in the 'long run', there is fair agreement on the circumstances in which the set of equilibria for an economy with $x$ times the money supply of an otherwise identical economy will be unchanged in real terms and will have all nominal prices multiplied by $x$. There is more controversy about the proper interpretation of cases in which purely nominal events have real consequences in some short run. The corresponding property of a growing economy is called 'superneutrality'.

Append to an otherwise routine neoclassical growth model a government that, whatever else it does, prints outside money and distributes it as transfers to the population. Each individual treats this transfer as disposable income along with wages and the interest or profits on real capital. There are now two vehicles for saving – cash balances and real capital – and households divide their savings between them according to some decision mechanism for portfolio balance. For equilibrium, all the real capital and all the money in existence must be willingly held by households. Suppose the government arranges for the money supply to grow at a constant rate $m$, and suppose the steady-state growth rate of output is $g$ (the sum of the growth rate of employment and the rate of labor-augmenting technological progress). Then in a steady state there will be inflation at the rate $m - g$. Now compare this steady state with another in which the money stock is growing at the rate $m'$ and inflation is at the rate $m'–g$ and, for definiteness, $m'$ exceeds $m$. Will output and capital stock be the same in the two economies? In that case the change in the money growth rate will have had no real effect in the long run; money is said to be superneutral.

In the particular model used by Tobin (1965), superneutrality does not hold. Instead, the higher inflation rate makes money a less desirable asset. This induces a shift in the economy-wide portfolio to hold more real capital and less money. In the new steady state, capital per worker is higher and therefore so is output per worker. Under the standard neoclassical institutional assumptions, greater capital intensity implies a higher steady-state real wage and a lower real interest rate. In that model, a change in the money growth rate has real effects of a systematic kind.

That is not the end of the story. In Tobin's version (1965), households allocate wealth between capital and real cash balances in a common-sense way depending on the nominal interest rate. Soon Sidrauski (1967) analyzed an otherwise similar model in which the representative household is an intertemporal utility maximizer who derives current utility from holding (real) money. In this model, money is superneutral, as one might expect. In steady state the real money stock grows as fast as real output, irrespective of the rate of inflation; and the path of the money supply has no effect on the productivity of real resources. So the optimizing household makes the same real intertemporal choices irrespective of the growth rate of nominal money.

That remark opens up further possibilities. On the assumption that money enters the economy's production function (i.e., that it allows a more efficient pattern of transactions), it is found that superneutrality is violated, as in the Tobin (1965) case but in the opposite direction. Faster money growth can lead eventually to lower steady-state capital intensity, because money is substituted for capital as an input. In that case even steady-state productivity can be lower.

Evidently the relation between money and growth depends critically on the real role that money plays in the economy. This branch of growth theory should probably not be thought to be predictive. Its role is to elucidate the relation between money and the real economy. One partial way to summarize the outcome is to say that the money-growth relationship turns on whether the rate of inflation has any long-run effect on the real interest rate. That is unlikely to have a simple institution-free answer.

## Income distribution

Very little has been said in this survey about income distribution (in other words, about the determination of factor prices). That is because there is no special connection between the neoclassical model of growth and the determination of factor prices. The usual practice is to appeal to the same view of factor pricing that characterizes static neoclassical equilibrium theory. If the working assumption that all markets clear were to be lifted, an alternative theory of factor prices would certainly be needed. Much else would change besides.

## REFERENCES

AGHION, P. and P. HOWITT (1992), "A model of growth through creative destruction", *Econometrica,* vol. 60, pp. 323-51.

ALLAIS, M. (1947), *Economie et intérêt,* Imprimerie National, Paris.

ARROW, K. (1962), "The economic implications of learning by doing", *Review of Economic Studies*, vol. 26, pp. 155-73.

ATKINSON, A. (1969), "The timescale of economic models: how long is the long run?", *Review of Economic Studies,* vol. 36, pp. 137-52.

AZARIADIS, C. (1993), *Intertemporal Macroeconomics,* Blackwell, Oxford.

AZARIADIS, C. and A. DRAZEN (1990), "Threshold externalities in economic development", *Quarterly Journal of Economics,* vol. 105, pp. 501-26.

BERTOLA, G. (1994), "Wages, profits, and theories of growth", in L. Pasinetti and R. Solow eds, *Economic Growth and the Structure of Long-Term Development,* St. Martin's Press, New York, pp. 90-108.

BLANCHARD, O. and S. FISCHER (1989), *Lectures in Macroeconomics,* The MIT Press, Cambridge, Mass.

BURMEISTER, E. (1980), *Capital Theory and Dynamics*, Cambridge University Press, Cambridge.

BURMEISTER, E. and A. DOBELL (1970), *Mathematical Theories of Economic Growth,* Macmillan, London.

CASS, D. (1965), "Optimum growth in an aggregative model of capital accumulation", *Review of Economic Studies,* vol. 32, pp. 233-40.

CASS, D. and K. SHELL (1983), "Do sunspots matter?", *Journal of Political Economy,* vol. 91, pp. 193-227.

DIAMOND, P. (1965), "National debt in a neoclassical growth model", *American Economic Review,* vol. 55, pp. 1126-50.

DIXIT, A. (1990), "Growth theory after thirty years", in P. Diamond ed., *Growth, Productivity, Employment,* The MIT Press, Cambridge, Mass., pp. 3-22.

DOMAR, E. (1946), "Capital expansion, rate of growth and employment", *Econometrica,* vol. 14, pp. 137-47.

FARMER, R. (1993), *The Macroeconomics of Self-Fulfilling Prophecies,* The MIT Press, Cambridge, Mass.

FELLNER, W. (1961), "Two propositions in the theory of induced innovations", *Economic Journal,* vol. 71, pp. 305-08.

FOLEY, D. and M. SIDRAUSKI (1971), *Monetary and Fiscal Policy in a Growing Economy,* Macmillan, New York.

GALE, D. (1967), "On optimal development in a multi-sector economy", *Review of Economic Studies,* vol. 34, pp. 1-18.

GROSSMAN, G. and E. HELPMAN (1991), *Innovation and Growth in the Global Economy,* The MIT Press, Cambridge, Mass.

HAHN, F. (1966), "Equilibrium dynamics with heterogeneous capital goods", *Quarterly Journal of Economics,* vol. 80, pp. 633-46.

HARROD, R. (1939), "An essay in dynamic theory", *Economic Journal*, vol. 49, pp. 14-33.

HELPMAN, E. (1992), "Endogenous macroeconomic growth theory", *European Economic Review,* vol. 36, pp. 237-67.

JOHANSEN, L. (1967), "A classical model of economic growth", in C. Feinstein ed., *Socialism, Capitalism, and Economic Growth*, Cambridge University Press, Cambridge, pp. 13-29.

JOHANSEN, L. (1974), *A Multi-Sectoral Study of Economic Growth*, 2nd enlarged edition, North Holland, Amsterdam.

KALDOR, N. (1961), "Capital accumulation and economic growth", in V. Lutz and D.C. Hagues eds, *The Theory of Capital,* Macmillan, London, pp. 197-222.

KENNEDY, C. (1964), "Induced bias in innovation and the theory of distribution", *Economic Journal,* vol. 74, pp. 541-47.

KING, R. and S. REBELO (1993), "Transitional dynamics and economic growth in the neoclassical model", *American Economic Review*, vol. 83, pp. 908-31.

KOOPMANS, T. (1965), "On the concept of optimal economic growth", *Scientific Papers of Tjalling C. Koopmans,* Springer, New York, pp. 485-547.

LEWIS, W.A. (1954), "Economic development with unlimited supplies of labour", *Manchester School*, vol. 22, pp. 139-91.

LOMBARDINI, S. (1994), "Growth and development", manuscript.

LUCAS, R. (1988), "On the mechanics of economic development", *Journal of Monetary Economics*, vol. 22, pp. 3-42.

MAHALANOBIS, P.L. (1953), "Some observations on the process of growth of national income", *Sankhya*, vol. 12, pp. 307-12.

MANKIW, G., D. ROMER and D. WEIL (1992), "A contribution to the empirics of economic growth", *Quarterly Journal of Economics*, vol. 107, pp. 407-48.

NELSON, R. (1981), "Research on productivity growth and productivity difference: dead ends and new departures", *Journal of Economic Literature*, vol. 19, pp. 2029-64.

ORPHANIDES, A. and R. SOLOW (1990), "Money, inflation and growth", in B.M. Friedman and F.H. Hahn eds, *Handbook of Monetary Economics*, vol. I, North Holland, Amsterdam, pp. 223-61.

PASINETTI, L.L. (1981), *Structural Changes and Economic Growth. A Theoretical Essay on the Dynamics of the Wealth of Nations,* Cambridge University Press, Cambridge.

PHELPS, E. (1961), "The golden rule of economic growth", *American Economic Review,* vol. 51, pp. 638-43.

RAMSEY, F. (1978), "A mathematical theory of saving", *Economic Journal*, vol. 88, pp. 543-99.

ROMER, P. (1986), "Increasing returns and long-run growth", *Journal of Political Economy*, vol. 94, pp. 1002-37.

ROMER, P. (1990), "Endogenous technological change", *Journal of Political Economy*, vol. 98, pp. S71-S102.

ROSENBERG, N. (1982), *Inside the Black Box: Technology and Economics,* Cambridge University Press, Cambridge.

SAMUELSON, P. (1958), "An exact consumption-loan model of interest, with or without the social contrivance of money", *Journal of Political Economy*, vol. 66, pp. 467-82.

SAMUELSON, P. (1965), "A theory of induced innovation along Kennedy-Weizsäcker lines", *Review of Economics and Statistics*, vol. 47, pp. 343-56.

SAMUELSON, P. (1975), "Optimum social security in a life-cycle growth model", *International Economic Rewiev*, vol. 16, pp. 539-44.

SIDRAUSKI, M. (1967), "Rational choice and patterns of growth in a monetary economy", *American Economic Review, Papers and Proceedings*, vol. 57, pp. 534-44.

SOLOW, R. (1956), "A contribution to the theory of economic growth", *Quarterly Journal of Economics*, vol. 70, pp. 65-94.

STIGLITZ, J.E. (1974), "Growth with exhaustible resources: the competitive economy", *Review of Economic Studies*, Symposium issue, pp. 123-38.

SWAN, T. (1956), "Economic growth and capital accumulation", *Economic Record*, vol. 32, pp. 343-61.

TOBIN, J. (1955), "A dynamic aggregative model", *Journal of Political Economy*, vol. 63, pp. 103-15.

TOBIN, J. (1965), "Money and economic growth", *Econometrica*, vol. 33, pp. 671-84.

UZAWA, H. (1961), "On the two-sector model of economic growth", *Review of Economic Studies*, vol. 28, pp. 40-47.

UZAWA, H. (1963), "On the two-sector model of economic growth", *Review of Economic Studies*, vol. 30, pp. 105-18.

VON WEIZSÄCKER, C. (1966), "Tentative notes on a two-sector model with induced technical progress", *Review of Economic Studies*, vol. 33, pp. 245-52.

YOUNG, A. (1993), "Invention and bounded learning by doing", *Journal of Political Economy*, vol. 101, pp. 443-72.