Investment-specific technological progress in the United Kingdom

*Hasan Bakhshi**
and
Jens Larsen†

* Structural Economic Analysis Division, Bank of England.
† Monetary Assessment and Strategy Division, Bank of England.

The views expressed are those of the authors and do not necessarily reflect those of the Bank of England. Thanks are due to John Butler for his help in the formative stages of this paper. We are grateful to Charlie Bean, Ian Bond, Tomasso Cosimano, Roy Cromb, Jo Cutler, John Duca, Simon Gilchrist, Andreas Hornstein, DeAnne Julius, Tiff Macklem, Steve Millard, Steve Nickell, Katharine Neiss, Nick Oulton, Mike Pakko, Peter Sinclair, Alison Stuart, and seminar participants at the BIS workshop on 'Structural change and inflation' in October 2000 for comments on earlier drafts of this paper. Thanks also to Morten Ravn for letting us use some of his software and to Nick Oulton for help with constructing and interpreting the data and the results.

Copies of working papers may be obtained from Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH; telephone 020-7601 4030, fax 020-7601 3298, e-mail mapublications@bankofengland.co.uk.

Working papers are also available at www.bankofengland.co.uk/workingpapers/index.htm

The Bank of England's working paper series is externally refereed.

Bank of England 2001
ISSN 1368-5562

Contents

Abstract

This paper analyses the impact of rapid technological change in the information and communications technology (ICT) sector on economic growth. Technological change in the ICT sector leads to a fall in the relative price of ICT goods, which leads firms to invest more heavily in high-tech goods. The approach is to build a dynamic general equilibrium model that is consistent with key stylised facts of the UK economy. We use that model to quantify the contribution to long-run growth of technological progress that is specific to the ICT sector. We find that technological progress that is specific to the ICT sector might account for around 20%-30% of long-run labour productivity growth, if this progress continues at the rate observed over the past 25 years. But this conclusion depends crucially on how ICT prices are measured and in particular on the estimate of the long-run rate of decline of ICT prices. We show that shocks to technological progress that is specific to production of ICT investment goods can have very different macroeconomic implications from a shock that applies to production of all goods. We demonstrate that a permanent increase in the growth rate of ICT-specific technological progress will increase the investment expenditure share of GDP but lower the aggregate depreciation rate, while an increase in the return to investment in ICT will increase both the expenditure share and the depreciation rate.

Journal of Economic Literature classification: E22, E32.

Key words: Investment-specific technological progress, TFP, ICT, growth accounting.

Summary

This paper addresses how rapid technological progress in the information and communications technology (ICT) goods sector contributes to long-run growth in the United Kingdom and how changes in the processes driving this progress may affect the macroeconomic outlook. Academics and policy-makers alike have argued that the US economy experienced an improvement in trend productivity growth in the second half of the 1990s. And technological progress in the ICT sector, with an associated rapid fall in ICT prices, has been a major contributor to US labour productivity growth over this period.

We start from the striking observation that the relative price of ICT goods has been declining steadily but at a very high rate over the past 23 years in the United Kingdom, accompanied by an increase in the real ICT investment to GDP ratio. We develop a model of the UK economy that can account for this fact. The mechanism is simple and intuitive: technological progress in the sector producing ICT capital goods leads to a decline in the relative price of ICT. Firms respond by substituting ICT capital for other types of capital and labour, raising the ICT capital intensity of production. In other words, technological progress in production of ICT capital contributes to output growth through relative price falls that induce capital deepening. In the model, we describe technological progress that applies only to the production of capital goods as investment-specific technological progress, and distinguish it from sector-neutral technological progress which applies to the production of all goods. The main difference between the two forms of progress is that investment-specific technological progress requires that investment is undertaken before it affects final output; sector-neutral technological progress is a 'free lunch' in that it affects final output directly.

Our model can be shown to be a special case of a more general framework, and has some additional appealing features. In particular, we can characterise the balanced growth path of our model of the UK economy, and can quantify the contributions that ICT investment-specific technological progress makes to long-run growth. This long-run growth path has the property that the expenditure share of ICT investment in GDP is constant: while the relative price of ICT is falling, the quantity of ICT relative to output increases, so the value of ICT investment relative to output stays constant. Our results suggest that despite ICT being only a relatively small component of the overall capital stock, ICT investment-specific technological progress contributes very significantly to labour productivity growth along the balanced growth path of our model of the UK economy, accounting for

around 20%-30% of labour productivity growth. But this conclusion depends crucially on how ICT prices are measured and the assumed rate of ICT price decline along the balanced growth path.

The paper goes on to consider various scenarios for structural change: first, if the rate of technological progress in ICT production increases temporarily, resulting in a temporary pick-up in the rate at which ICT prices decline; second, if the rate of technological progress increases permanently, and third, if structural changes lead to temporary increases in the expenditure share of ICT investment in overall output. We show that this last scenario can account for the increase in the rate at which the aggregate capital stock depreciates, as appears to have been observed in the United Kingdom in the 1990s.

1 Introduction

A broad consensus appears to have emerged amongst academics and policy-makers alike that there has been some improvement in (at least medium-term) US trend productivity growth in the second half of the 1990s. Recent attempts to decompose US labour productivity growth into its main determinants find that information and communications technology (ICT) has made significant contributions through increases in both capital deepening and total factor productivity (TFP) growth over this period. Notable examples include the work by Oliner and Sichel (2000), Gordon (2000), Jorgenson and Stiroh (2000) and Whelan (2000a). Kneller and Young (2000) and Oulton (2000) perform similar decompositions for the United Kingdom, though the data constraints are greater in this case.(1)

In this paper, this approach is labelled 'historical growth accounting'. A separate literature using dynamic general equilibrium (DGE) models distinguishes between technological progress that is specific to production of capital goods and technological progress that is 'neutral' in the sense that it applies to production of all goods (TFP).(2) The main reference here is Greenwood, Hercowitz and Krusell (1997), with Greenwood, Hercowitz and Krusell (2000) and Pakko (2000) being recent examples of applications. These models do not attempt historical decompositions of labour productivity growth, but instead decompose productivity along the balanced growth path of the economy into investment-specific technological progress and neutral technological progress. This approach emphasises the importance of substitution effects: rapid technological progress in production of capital goods leads to declining prices and hence to increasing capital intensity.

In this paper, we adapt the model of Greenwood, Hercowitz and Krusell (1997) to quantify the contribution of ICT-specific technological progress to productivity growth along the balanced growth path of a model of the UK economy, drawing heavily on the efforts of our colleague Nick Oulton at the Bank of England to derive ICT investment data for the United Kingdom. Like Greenwood, Hercowitz and Krusell (1997), our motivation is the observation that rapid declines in the relative

(1) One of the issues for these papers is that it is not at all clear that there has yet been any increase in trend productivity growth in the United Kingdom, despite strong increases in ICT investment. That is not of course to rule out the possibility that such productivity improvements might be on the horizon.

(2) This relates directly to the Solow (1960) and Jorgenson (1966) debate on whether technological progress is 'embodied' or 'dis-embodied'. Hercowitz (1998) argues that this language is imprecise and instead uses the distinction between 'sector-specific' and 'neutral' technological progress that we also use in this paper.

Figure 1: Relative price of ICT goods and the ICT-output ratio

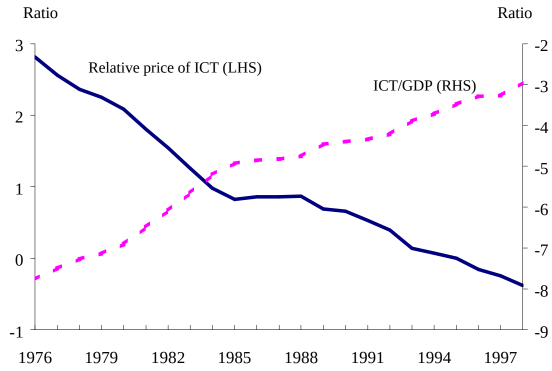

Notes: (1) All series are in logs. (2) ICT is measured in real quantities.

price of ICT goods have been accompanied by an increase in the ratio of real ICT investment, measured in units of ICT, to (non-housing) output (see Figure 1).[(3)] We identify technological progress in production of ICT goods as inversely related to the relative price of ICT goods. Using this information and the model's balanced growth path relations, we can calculate the contribution of ICT-specific technological progress to labour productivity growth along the balanced growth path. We find that despite the fact that ICT is a relatively small component of the overall capital stock, ICT-specific technological progress contributes significantly to labour productivity growth along the balanced growth path for the UK economy, accounting for around 20%-30% of labour productivity growth.

The key advantage of the DGE approach over growth accounting is that it permits forward-looking analysis: the short-run macroeconomic implications of a shock to investment-specific technological progress or TFP can be simulated, even if such shocks have not yet hit the economy. This is a particularly useful tool in our context, as it provides a macroeconomic guide for policy-makers who wish to incorporate such shocks into their forecasts. We present impulse responses for temporary shocks to both ICT-specific technological progress and to neutral technological progress. Shocks to ICT-specific technological progress have very different implications for investment, depreciation, the capital stock and labour productivity than shocks to neutral technological progress. The main driver of these differences is that where an increase in sector-neutral technological progress has an immediate 'free-lunch' effect on final output—final output increases for a given level of factor inputs—technical progress that is specific to production of ICT investment

[(3)]The details of how these series are derived are discussed at length in Section 3.

goods requires that investment is undertaken. We describe these effects using a simple baseline model, but also consider extensions and variations that arguably bring the model closer into line with certain empirical regularities. In particular, we consider modifications to the labour supply specification, capital adjustment costs and variable utilisation of capital, and also modify the specification of the stochastic processes driving the shocks.

The main disadvantage of this approach is that it necessarily loses some of the empirical richness of the historical growth accounts. In particular, the balanced growth decompositions of Greenwood, Hercowitz and Krusell (1997) ignore the contribution that ICT makes to labour productivity through the direct effect of TFP improvements in the ICT-producing sector on economy-wide TFP. As such this would understate ICT's contribution to long-run growth. Against this though, Hercowitz (1998) notes that the treatment of investment-specific technological progress in Greenwood, Hercowitz and Krusell (1997) implicitly assumes there are no resource costs to the economy when enjoying investment-specific technological progress. This is likely to overstate the contribution of ICT to long-run economic growth. In the following, we spell out in more detail the relation between the two approaches to growth accounting.

1.1 Balanced growth and 'historical' growth accounting

The balanced growth accounting exercise differs from 'historical' growth accounting by focusing on the long-run, or steady-state, growth path. Growth accounting is about attributing growth at a particular point in time to growth in factor inputs and total factor productivity, taking prices and quantities as given. Take a typical but stylised growth accounting equation:

$$\Delta \ln Y_t = \alpha_t \Delta \ln K_t + (1 - \alpha_t)(\Delta \ln N_t + \Delta \ln H_t) + \Delta \ln TFP_t, \text{ or} \qquad (\mathbf{1})$$

$$\Delta \ln Y_t - (\Delta \ln N_t + \Delta \ln H_t) = \alpha_t (\Delta \ln K_t - \Delta \ln N_t - \Delta \ln H_t) + \Delta \ln TFP_t \qquad (\mathbf{2})$$

In (**1**), output growth $\Delta \ln Y_t$ is attributed to growth in capital inputs $\Delta \ln K_t$, labour inputs in heads and hours $(\Delta \ln N_t + \Delta \ln H_t)$, weighted by their (possibly time-varying) income shares, and to growth in total factor productivity, $\Delta \ln TFP_t$. (**2**) is a simple rearrangement that attributes growth in labour productivity, measured per hour, to capital deepening, ie an increase in the capital-labour ratio, and to total factor productivity. Ignoring statistical issues, this is an accounting identity: indeed total factor productivity growth is calculated to make these equations hold with equality.

These equations are obviously useful tools for providing a historical account of output or productivity growth. But they are less useful as a tool for forward-looking analysis: by taking factor inputs as given, growth accounting does not provide us with a tool for making projections for future growth, because it is conditional on the behaviour of factor demand. The DGE approach differs by characterising a steady-state balanced growth path of a dynamic general equilibrium model that imposes constraints on factor inputs. Specifically, the steady-state balanced growth path is characterised by constant growth rates. Growth in capital inputs is related to growth in its economic determinants, neutral technological progress and to investment-specific technological progress. Employment grows at a constant rate, ie the rate of population growth and hours per worker are constant. Income shares are constant. In other words, along the balanced growth path:

$$\Delta \ln Y = \alpha (\Delta \ln Y + \Delta \ln Q) + (1 - \alpha)(\Delta \ln N) + \Delta \ln TFP \qquad \textbf{(3)}$$

where no subscript indicates that the variable is time-invariant. Here, capital growth is characterised as the growth in production of final goods $\Delta \ln Y$ (as this is a homogeneous good model) and the growth that is specific to production of investment goods, $\Delta \ln Q$. This equation is useful because, unlike (**1**) and (**2**), it characterises the long run.[(4)] In the following we describe Q as sector-specific while TFP is described as sector-neutral technological progress; notation-wise, we use the term Z to describe TFP.

Greenwood, Hercowitz and Krusell (1997) offer two alternative interpretations of the index, Q. First, in this homogeneous good model, Q can be seen as denoting the amount of capital that can be purchased in efficiency units for one unit of final output. This increases over time with investment-specific technological progress. A second interpretation is that Q represents the vintage of a capital good: each period a new vintage is produced that is successively more productive—of 'higher quality'—than the previous one. The empirical counterpart of Q is identical in both interpretations: it equals the inverse of the price of investment goods, adjusted for quality, relative to some measure of the price of the homogeneous good. In our model, this must be a consumption deflator as the homogeneous good enters agents' utility functions.

[(4)] As mentioned above, the disadvantage of this framework is that it is necessarily less rich than a growth accounting framework à la Jorgenson: in this example, and in our balanced growth accounting, we do not take account of factors such as labour quality that are obviously important in providing an account of economic growth.

In the growth accounting literature, the expenditure measure of GDP growth includes a measure of investment that allows for the 'quality' of capital goods having improved over time. The empirical implication is that the left-hand side of (**3**) should be deflated by a quality-adjusted deflator to reflect the quality improvement in the investment component of aggregate demand. In the homogeneous good framework of Greenwood, Hercowitz and Krusell (1997), no such allowance is made. In this literature, output is expressed in units of the homogeneous good and so the empirical counterpart is that output should be deflated by a consumption deflator.[(5)] Hercowitz (1998) sets out a framework that he argues nests the positions of both these traditions. In particular, he shows that the homogenous good model embedded in Greenwood, Hercowitz and Krusell (1997) assumes there are no resource costs to the economy from investment-specific technological progress, while arguing, following Hulten (1992), that quality-adjusting the left-hand side of (**3**) is a way of incorporating such resource costs: an increase in quality requires a reduction in some other expenditure component for a given level of aggregate output. In a one-sector model, this has undesirable implications; in particular, the relative price of investment goods is constant, inconsistent with the empirical evidence, and the difference between investment-specific and sector-neutral progress can no longer be identified. Hercowitz's (1998) essay implies that a more general model that allows for *some* form of resource cost would be superior. In the absence of such a model, we follow Greenwood, Hercowitz and Krusell (1997, 2000), implicitly assuming that there are no resource costs of investment-specific technological progress.[(6)]

The remainder of the paper is organised as follows. In Section 2 we set out the baseline model, characterising the equilibrium of the model and its balanced growth path. In Section 3 we calibrate the baseline model to the UK economy, and decompose labour productivity growth along the balanced growth path into investment-specific technological progress and neutral technological progress. Section 4 presents the dynamic analysis of the baseline model, drawing out the key differences in the macroeconomic effects of investment-specific shocks and neutral

(5) This assumes that the economy is closed. In an open economy the empirical counterpart should, strictly speaking, be the domestically produced component of the consumption deflator. This is consistent with the homogeneous good assumption if we assume that countries specialise in production and that all imports are final goods.

(6) Jorgenson and Stiroh (2000) further argue that the investment-specific technological progress identified by Greenwood, Hercowitz and Krusell (1997) as accounting for the major component of post-war US economic growth in fact reflects disembodied technological progress in the production of semi-conductors used as intermediate inputs. We do not comment on this, except noting that in our homogeneous good model such a distinction cannot be made.

shocks to technological progress. Section 5 presents extensions of the baseline model. We choose those extensions from the existing theoretical literature as these address some obvious shortcomings of the baseline model. Section 6 considers some 'scenarios for structural change': more specifically, we consider the dynamic implications of permanent rather than temporary shocks to the level of technology, and draw out some implications of changing the growth rate of technological progress and the return to investment in some comparative statistics exercises. Finally, Section 7 concludes.

2 The baseline model

In the following, we describe the baseline model and characterise equilibrium and the balanced growth path. The model follows Greenwood, Hercowitz and Krusell (1997) closely, with the main differences being that we split the capital stock into ICT (indexed by e for exciting) and non-ICT (indexed d for dull) capital rather than equipment and structures, and we allow for investment-specific technological growth in ICT.(7) This latter distinction makes the analysis more relevant to the current UK policy debate.

The key characteristic that distinguishes this model from a standard one-sector growth model is the capital accumulation equation. In the current model, the stock

(7) The non-stationarity of the quality-adjusted equipment investment to GDP ratio and the stationarity of the structures investment to GDP ratio in the United States is used by Greenwood, Hercowitz and Krusell (1997) to motivate their assumption that there is sector-specific technological progress in equipment but not in structures. But the structures investment data are not quality-adjusted in the same way as the equipment data in the United States: certainly no hedonic adjustments are made (although the Bureau of Labour Statistics does make other adjustments). Even with sector-specific technological progress in structures, the ratio of non-quality adjusted structures to GDP would be stationary along the balanced growth path. So this is not in fact a good motivation for their assumption of no sector-specific technological progress in structures. Our series for the quality-adjusted ICT ratio with respect to GDP is non-stationary in the United Kingdom too, so we assume that investment-specific technological progress occurs in that sector. As with the structures data in the United States, however, our non-ICT data are not hedonically adjusted, even though rapid quality improvements are likely to have occurred outside the ICT sector (see Gordon (1990)). In the absence of such data for non-ICT investment, we do not allow for sector-specific technological progress in the non-ICT sector. This characterisation of no sector-specific technological progress in the non-ICT sector is not therefore directly comparable with Greenwood, Hercowitz and Krusell's analysis for the United States insofar as non-ICT investment contains non-ICT elements of equipment that Greenwood, Hercowitz and Krusell assume is subject to sector-specific technological progress. Gort, Greenwood and Rupert (1999) use a panel dataset on rental values to estimate sector-specific technological progress in US structures investment too.

of capital of type $i = d, e$ at time $t+1$, K^i_{t+1} is related to the stock of capital and investment at time t, X^i_t, through:

$$K^i_{t+1} = (1 - \delta_i)K^i_t + Q^i_t X^i_t \tag{4}$$

where $Q^d_t \equiv 1$ and δ_i is a parametric depreciation rate.[8]

The factor Q^i_t determines the amount of capital of type i that can be purchased for one unit of final output; in the standard neoclassical growth model $Q^i_t \equiv 1$ but here we allow Q^e_t to increase over time. Notice that investment X^i_t is measured in units of final goods, so aggregate investment X_t is given by $X_t = \sum X^i_t$. Here, we interpret Q^i_t as a measure of technological change specific to the production of investment good i: a rise in Q^i_t lowers the marginal cost of producing investment goods measured in units of final goods, and so Q^i_t is inversely related to the relative price of capital good i. One simple way to spell out this relationship and to outline the sectoral interpretation of the model is the following: ICT capital goods are produced by firms, using materials M^i_t as the only input in the production process, charging a price P^i_t for their output in a perfectly competitive market. Such a firm maximises profits $P^i_t\,(Q^i_t M^i_t) - M^i_t$ where $(Q^i_t M^i_t)$ is the firm's output and the price of materials, in the form of final goods, is normalised at one. The first-order condition for this problem, where the firm determines its output levels taking prices and technology as given, together with a zero profit constraint implies that $P^i_t = 1/Q^i_t$. We use this relationship in the calibration exercise, where the growth rate of Q^i_t is calibrated using series on relative prices of ICT capital goods.

In the model below, we follow a convention whereby capital letters denote trended variables and lower-case letters indicate stationary variables. All *quantity* variables are measured in *per capita* terms.

2.1 The agents

The economy is inhabited by an infinitely lived, representative agent or household who has time-separable preferences U defined over consumption C_t of final goods and leisure L_t. The agent chooses C_t, L_t and investment X_t to maximise the expected present value of contemporaneous utility, using a discount factor β, subject to a budget constraint:

[8] Fraumeni (1997) reports that geometric depreciation is in general a good approximation to the decline of asset prices with age.

$$\max E \sum \beta^t U\left(C_t, 1-h_t\right)$$

$$s.t.\;\; C_t + X_t = (1-\tau_k)(r_t^d K_t^d + r_t^e K_t^e) + (1-\tau_l) W_t h_t + T_t. \qquad (\mathbf{5})$$

Here, consumption and investment cannot exceed the sum of labour and capital rental income net of taxes and lump sum transfer, T_t; wages and hours worked are W_t and h_t respectively, and τ_l is the tax rate on labour income. Rental income has two components: there is rental income from capital of type d at rate r_t^d and type e at rate r_t^e, with quantities at K_t^d and K_t^e respectively. The tax rate on rental income is τ_k.

The agent's capital holdings of type $i = d, e$ evolve according to (**4**) reported below as (**6**), for convenience:

$$\gamma^L K_{t+1}^i = (1-\delta_i) K_t^i + Q_t^i X_t^i, \qquad (\mathbf{6})$$

where $Q_t^d \equiv 1$ at all times, δ_i is the depreciation rate for capital of type i and γ^L is the deterministic growth rate of the population or, equivalently, of the household.

The agent maximises expected lifetime utility subject to the budget constraint (**5**), and the accumulation equations (**6**) by choosing C_t, L_t and X_t^i. The first-order conditions for this problem are:

$$\begin{aligned} U_l(C_t, 1-h_t) &= U_c\left(C_t, 1-h_t\right)\left(1-\tau_l\right) W_t \\ \lambda_t^e &= \lambda_t^d = U_c\left(C_t, 1-h_t\right) \\ \lambda_t^i \gamma^L / Q_t^i &= \beta E_t \lambda_{t+1}^i \left((1-\tau_k) r_{t+1}^i + \beta(1-\delta_i)/Q_{t+1}^i\right),\; i = d, e \end{aligned} \qquad (\mathbf{7})$$

where subscripts c and l indicate derivatives of $U()$ with respect to the first and second argument respectively. The first condition equates the marginal disutility from an additional hour of work with the marginal return to working, adjusted for taxes and measured in utility terms. The second condition describes the marginal utility of an additional unit of capital of type i: as there are no additional resource costs associated with changing capital from one type to the other, the marginal utilities of an additional unit of the capital goods are equal. And as capital goods can be transformed into consumption goods at no cost, the marginal utility of an additional unit of capital equals the marginal utility of consumption. The third condition is the standard Euler equation, equating the marginal cost of acquiring an additional unit of capital today in utility terms with the discounted expected return to this investment, consisting of expected after-tax rental income and the value of having this unit next period, adjusted for depreciation and possible capital losses.

2.2 Firms

In the baseline model, Q_t^i is assumed to capture all differences between production of final and investment goods: apart from technological progress, the production process is identical across goods. So a characterisation of firms producing final goods is sufficient. The firms in this economy have access to a production technology for final goods that uses capital of both types and labour:

$$Y_t = F(K_t^e, K_t^d, Z_t h_t) \tag{8}$$

where Y_t is output and Z_t is labour-augmenting technological progress that applies to production of all goods. F is assumed to be continuous and concave in each of the inputs, and homogenous of order one. Goods and factor markets are assumed to be perfectly competitive, so that firms in their production decisions take output and factor prices as given. Firms rent capital and labour on a period-by-period basis—the workers hold the capital stock—so the firms' dynamic optimisation problem is identical to a sequence of the following static optimisation problems:

$$\max \ \Pi_t = F(K_t^e, K_t^d, Z_t h_t) - r_t^e K_t^e - r_t^d K_t^d - W_t h_t \tag{9}$$

The first-order conditions for this problem are:

$$\begin{aligned} F_{K^i}(K_t^e, K_t^d, Z_t h_t) &= r_t^i;\ i = d, e && (10)\\ Z_t F_h(K_t^e, K_t^d, Z_t h_t) &= W_t && (11) \end{aligned}$$

where subscript $K^i, i = d, e$ and h indicate derivatives with respect to the first, second and third argument. There are no dynamic aspects to the firms' decisions, so the conditions describing factor demand simply state that marginal cost, given by real rental rates and real wages, equal marginal factor products, given by the marginal products of capital and labour.

2.3 Government

We incorporate a tax-levying government in the model because of the potentially important effects that distortionary taxes have on capital accumulation, and hence on the decomposition exercise. We are not analysing the use of taxation in demand management in this paper, and simply assume that the government balances its budget period-by-period by returning revenues from distortionary taxes to the agents via lump-sum transfer. The government's budget constraint is then:

$$T_t = \tau_k(r_t^e K_t^e + r_t^s K_t^s) + \tau_l W_t h_t \tag{12}$$

This completes the description of the baseline model. In the following, we characterise equilibrium and the balanced growth path.

2.4 Equilibrium and balanced growth

To facilitate our exposition of the steady state, we make assumptions about particular functional forms here. We assume a Cobb-Douglas production function[(9)] and a logarithmic specification for the instantaneous utility function:

$$\begin{aligned} Y_t &= Z_t \left(K_t^e\right)^{\alpha_e} \left(K_t^d\right)^{\alpha_d} \left(h_t\right)^{1-\alpha_e-\alpha_d} && \textbf{(13)} \\ U\left(C_t, L_t\right) &= \theta \ln\left(C_t\right) + (1-\theta)\ln\left(1-h_t\right) && \textbf{(14)} \end{aligned}$$

Prior to characterising the balanced growth path, we describe the equilibrium of this economy. Equilibrium is characterised by a set of time-invariant decision rules for C_t, X_t^i and h_t, pricing functions for W_t, r_t^i, a balanced budget rule, and laws of motion for the aggregate capital stock that solve the agents' and firms' optimisation problem and satisfy the economy's resource constraint. These conditions are summarised by the following set of equations:

$$\begin{aligned} \frac{h_t}{(1-h_t)} &= \frac{\theta}{1-\theta}(1-\tau_l)(1-\alpha_e-\alpha_d)\frac{Y_t}{C_t} && \textbf{(15)} \\ \lambda_t^i \gamma^L / Q_t^i &= \beta E_t \lambda_{t+1}^i \left((1-\tau_k)\alpha_i \frac{Y_{t+1}}{K_{t+1}^i} + (1-\delta_i)/Q_{t+1}^i\right),\ i=d,e && \textbf{(16)} \\ \lambda_t^e &= \lambda_t^d = \frac{\theta}{C_t} && \textbf{(17)} \\ \gamma^L K_{t+1}^i / Q_t^i &= (1-\delta_i)K_t^i/Q_t^i + X_t^i,\ i=d,e && \textbf{(18)} \\ C_t + X_t^e + X_t^d &= Z_t \left(K_t^e\right)^{\alpha_e} \left(K_t^d\right)^{\alpha_d} \left(h_t\right)^{1-\alpha_e-\alpha_d} && \textbf{(19)} \end{aligned}$$

The first three conditions, (**15**)–(**17**), come straightforwardly from combining the first-order conditions characterising the agents' problem with those characterising the firms' and hence need no further comment. (**18**) characterises the economy's accumulation of capital of type i. The resource constraint, (**19**), is obtained from combining the budget constraint of the worker with the government budget constraint, using the homogeneity properties of the production function.

We can now characterise the non-stochastic, steady-state balanced growth path of this model as an equilibrium satisfying conditions (**15**)–(**19**) where all variables

[(9)]Notice that we have detached the technology variable Z_t from labour inputs here, so we have not written Z_t as labour-augmenting. This is purely for convenience and ease of comparison with Greenwood, Hercowitz and Krusell (1997): with a Cobb-Douglas production function, labour-augmenting and factor-neutral technological progress are identical up to a constant.

grow at a constant rate. Denote the *gross growth rate* of output per capita, Y_t, along the balanced growth path with g, and of capital per capita, K_t^i, with g_i.[10]

A balanced growth path obviously requires that hours per worker do not grow (otherwise they will eventually hit their upper or lower bound). Combined with the fact that this is a full-employment economy, this implies that total hours grow at the rate of population: the only contribution from hours worked to output growth comes from growth in labour force, and ultimately, as participation rates along a balanced growth path are constant, from population growth. In the model, we assume no population growth along the balanced growth path, to ease the description and facilitate comparison with Greenwood, Hercowitz and Krusell (1997), ie assuming that $\gamma^L = 1$. This has no implications for the growth accounting exercise, as we are accounting for labour productivity growth, which, by nature, is independent of the size of the population, but would obviously affect an estimate of the growth rate of aggregate output along a balanced growth path.

From (**19**), balanced growth requires that the demand components of the model, ie C_t, X_t^d and X_t^e, grow at the same gross rate as output Y_t, ie g. Furthermore, let γ_e and γ_z describe the steady-state gross growth rates of Q_t^e and Z_t. Using the production function, this implies that:

$$g = \gamma_z g_e^{\alpha_e} g_d^{\alpha_d} \qquad \textbf{(20)}$$

From (**20**), in the long run, increases in output can be accounted for by neutral technological progress or, equivalently because the production function is Cobb-Douglas, by labour-augmenting technological progress, γ_z, and by increases in the capital stock per capita, equivalent to capital deepening, $g_e^{\alpha_e}$ and $g_d^{\alpha_e}$. But growth in the capital stock depends on technological progress in production of ICT capital goods, *in addition to* neutral technological progress. The dependence stands out from the capital accumulation equations, where, by (**18**), $g_e = g\gamma_e$. Notice that the assumption that $Q_t^d \equiv 1$ implies that $g_d = g$. Combining this with (**20**), the growth rates can be expressed as functions of the exogenous growth rates of the production technologies:

$$\begin{aligned} g &= \gamma_z^{1/(1-\alpha_e-\alpha_d)} \gamma_e^{\alpha_e/(1-\alpha_e-\alpha_d)} \\ g_e &= \gamma_z^{1/(1-\alpha_e-\alpha_d)} \gamma_e^{(1-\alpha_d)/(1-\alpha_e-\alpha_d)} \end{aligned} \qquad \textbf{(21)}$$

The equilibrium conditions (**15**)–(**19**) can now be transformed by expressing them

[10] So, for example, a growth rate of 2% is a gross growth rate of 1.02.

in terms of the following variables, where lower case indicates stationary variables:

$$
\begin{aligned}
y_t &= Y_t/g^t;\ c_t = C_t/g^t;\ x_t^e = X_t^e/g^t;\ x_t^d = X_t^d/g^t;\ k_t^d = K_t^d/g^t;\ k_t^e = K_t^e/g_e^t \\
q_t^e &= Q_t^e/\gamma_e^t;\ z_t = Z_t/\gamma_z^t;\ \tilde{\lambda}_t^e = \lambda_t^e g^t;\ \tilde{\lambda}_t^d = \lambda_t^d g^t \qquad \textbf{(22)}
\end{aligned}
$$

These variables are stationary, so a balanced growth path with constant growth in the non-normalised variables can be characterised as a *stationary state* with no growth in these transformed variables. Let no time subscript indicate stationary-state values. Then the balanced growth path is characterised by the following set of equations:

$$
\begin{aligned}
\frac{1-\alpha_d-\alpha_e}{(1-\tau_l)}\frac{1-\theta}{\theta}\frac{y}{c} &= \frac{1-h}{h} && \textbf{(23)} \\
\frac{\beta}{g\gamma_i}\left((1-\tau_k)\alpha_i\frac{y}{k^i}+(1-\delta_i)\right) &= 1,\ i=d,e && \textbf{(24)} \\
\frac{k^i}{y}\left(g\gamma_i-(1-\delta_i)\right) &= \frac{x^i}{y},\ i=d,e && \textbf{(25)} \\
\tilde{\lambda}^e=\tilde{\lambda}^d &= \frac{\theta}{c} && \textbf{(26)} \\
\frac{c}{y}+\frac{x^e}{y}+\frac{x^d}{y} &= 1 && \textbf{(27)}
\end{aligned}
$$

Before moving on to assessing the importance of investment-specific technological progress in accounting for long-run growth, it is worth characterising the steady-state growth path in words. Along the steady-state path, productivity in the production of ICT capital goods is increasing faster than productivity in production of consumption goods, so the relative price of ICT capital is falling at a constant rate. The ICT capital-labour ratio is increasing faster than labour productivity, so ICT capital deepening is faster than output growth. Investment in ICT, measured in units of ICT goods, increases in line with the ICT capital stock and hence faster than output, but due to falling prices, ICT investment expenditure grows in line with output, so the investment expenditure share of GDP stays constant.

3 Characterising the balanced growth path

To assess the contribution of investment-specific technological progress to long-term growth, the parameters of the model must be assigned values. We follow the calibration approach advocated by Kydland and Prescott (1982). According to this approach, parameter values are set either according to related empirical evidence or, in the absence of such evidence, to ensure that the model's balanced growth path is consistent with averages observed in UK aggregate data over the sample period (1976-98). Consistency with the balanced growth path is an important feature of this approach—the parameter values must be set consistently such that for the chosen set of parameters, the equations characterising the balanced growth path, (**23**)—(**27**), are satisfied. In this sense, the model guides our interpretation of the data.

3.1 Calibration

The parameters of the model are

$$\{\theta, \beta, \alpha_e, \alpha_d, \delta_d, \delta_e, \gamma_e, \gamma_z, \tau_l, \tau_k\}.$$

The growth rate γ_e is calibrated directly using a deflator for ICT investment goods.(11) It is the growth rate of the hedonic investment deflator minus the growth rate of the consumption deflator.

Reliable hedonic deflators for ICT goods that attempt to control for quality improvements are not available in the United Kingdom. In the absence of such data, we follow Broadbent and Walton (2000), Kneller and Young (2000) and Oulton (2000) in employing a law of one price-type argument and use deflators from the United States, converted to £ using the $-£ exchange rate.(12)(13) In particular, we use estimates of nominal investment expenditure on computers, software and telecommunications in the United Kingdom derived from input-output tables (see Oulton (2000)(14)), to weight together computer, software and telecommunications deflators from the US National Income and Product Accounts.(15) We treat the resulting chain-weighted Fisher price indices as our ICT investment deflator series. This is expressed relative to the non-durable goods (excluding housing services)

(11) K_d, while representing 'dull' capital, is productive capital, so excludes housing capital. This is appropriate because our measure of output, Y_t, excludes housing services.

(12) Gust and Marquez (2000) discuss how Australia, Denmark and Sweden all officially use US hedonic computer deflators, exchange rate-adjusted, to proxy quality-adjusted computer prices in their respective countries.

(13) Because ICT products are traded on a global market, it seems likely that the rate at which quality-adjusted prices are falling over time should be the same in the United Kingdom and United States. The *level* of prices may differ, say because of market discrimination by suppliers who possess monopoly power. But even changes in the degree of monopoly power are likely to be swamped by the huge falls in US prices related to investment-specific technological progress.

(14) Oulton (2000) notes that while the growth rates of software investment in nominal terms has been similar in the United States and United Kingdom in the official data, the level of UK software relative to computers investment is much smaller in the United Kingdom. Oulton suggests that an upward adjustment be made to the UK data to control for this.

(15) There are currently no official data available in the United Kingdom for our definition of ICT investment.

Figure 2: Relative price of ICT goods and the ICT-output ratio. 'High software'

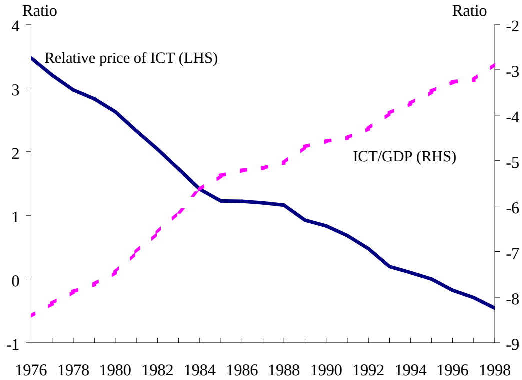

Notes: (1) All series are in logs. (2) ICT is measured in real quantities.

deflator. It is the average growth rate in the resulting series over our sample that gives us γ_e.

According to Moulton, Parker and Seskin (1999) and Parker and Grimm (2000), only prices for pre-packaged software in the US NIPA are calculated from constant-quality price deflators based on hedonic methods. Prices for firms' own-account software in the NIPA are based on input cost indices that implicitly assume no increase in the productivity of programmers. Custom software prices are assumed to be a weighted average of pre-packaged software prices and own-account software (with an arbitrary weight of 75% on own-account software). But it is implausible to assume that the productivity of programmers has not improved over time. This might lead to a significant understatement in the decline in the relative price of software and hence in our ICT deflator. To investigate the implications of this possible mis-measurement for assessing the importance of ICT investment-specific technological progress, we also present balanced growth accounting estimates calculated on the assumption that pre-packaged software prices capture price trends for all types of software (we refer to this variant as the 'high software' case as distinct from the 'low software' case consistent with NIPA data).(16) The 'high software' relative price and quantity-output ratio are reported in Figure 2 while the 'low software' is the data underlying Figure 1.

Of course, ICT goods are not the only types of investment good that have been

(16)Jorgenson and Stiroh (2000) go further still and report traditional growth accounting estimates under the assumption that software prices fall at the even more rapid rate reported by Brynjolfsson and Kemerer (1996) for microcomputer spreadsheets in 1987-92.

subject to quality improvement (see Gordon (1990)). Hedonic price measurement by the Bureau of Economic Analysis in the United States is restricted to ICT goods. So the data constraints we face in deriving a plausible quality-adjusted non-ICT deflator are particularly severe. Given this, we assume that there is no sector-specific technological progress in the production of non-ICT capital goods.

The growth rate g is calibrated by estimating average labour productivity growth over the sample, where the hours worked series is based on New Earnings Survey and Labour Force Survey data. The within-sample properties of hours per capita and labour force participation differ from those of a balanced growth path: it is well-known that since 1976, average hours per worker have declined and participation rates in the United Kingdom have increased. The correct way to estimate output/productivity growth along a balanced growth path where such changes are not possible is to control for these factors within sample: we hence measure output growth per hour, and infer the long-run output growth by combining this measure with the balanced growth requirement that hours per worker and participation rates are constant.

The depreciation parameters δ_d and δ_e are key parameters in the construction of the ICT and non-ICT capital stocks using (**6**). For δ_e, we use the time series for the current price capital stock of computers, software and telecommunications in Oulton (2000)(17) to chain-weight together the depreciation rates for computers, software and telecommunications in Jorgenson and Stiroh (2000).(18) The sample average (1976-98) of the resulting weighted average depreciation rate series is 0.241 assuming software low and 0.239 assuming software high. The depreciation rate for non-ICT capital, δ^d, is derived in a similar fashion, using estimates of the non-ICT capital stock taken from Oulton (2000) and depreciation rates for vehicles, plant and intangible fixed assets from Fraumeni (1997).(19) The sample average is 0.056.

(17)These capital stock series are chain-weighted series. Whelan (2000b) shows that in this case current price, not constant price, capital stocks are the appropriate weights when calculating a weighted average depreciation rate. Following Greenwood, Hercowitz and Krusell (1997), the measure used for the capital stock is a wealth measure. If a productive measure of the capital stock were used, for example an index of the volume of capital services, the appropriate weights to use in our calculation would be rental prices.

(18)Specifically, we assume depreciation rates of 31.5% per year for computers and software and 11% per year for telecommunications.

(19)This method of calculating ICT and non-ICT capital stocks produces estimates of the real wealth stock at replacement value. The economic depreciation rates, δ_t^j, denote the decline in the replacement value of a unit of capital (relative to the price of new capital) that occurs as the unit ages. But it is the real productive capital stock that enters into the production function in (**13**).

With these parameters determined, the balanced growth path investment-capital ratios can be determined from the capital accumulation equations (**25**). We then measure the ratios x^i/y using the data from Oulton (2000). Given that we use the same deflator for both investment and output, these can be measured in nominal or real terms. From these we can infer the consumption–output ratio $c/y = 1 - \Sigma x^i/y$.[20]

From the income side of National Accounts, a steady-state labour share of 70% is estimated. A marginal tax rate on labour income of 42.7% is used, τ_l, based on the work by Millard, Scott and Sensier (1997). This is the average value of the marginal tax rate faced by a worker on average earnings over the period 1976-98. Specifically, it is the basic rate of income tax plus the marginal national insurance contribution faced by such a worker, divided by one plus the marginal national insurance contribution faced by their employer. With Cobb-Douglas and perfect competition, the labour share is equal to $1 - \alpha_d - \alpha_e$. We also determine h, the proportion of hours available used for work as 0.26. This is the average portion of non-sleeping time spent in work reported in two 'use of time' studies in the United Kingdom discussed by Jenkins and O'Leary (1997). This is very similar to the 0.24 used by Greenwood, Hercowitz and Krusell (1997) for the United States. With these estimates at hand, the first-order condition for labour characterising the balanced growth path, (**23**), determines the utility parameter θ.

So the appropriate depreciation rate is actually a physical decay rate: the rate at which a unit of capital of a given vintage becomes less capable of producing output as it ages. In a simple model of vintage capital with investment-specific technological progress, Whelan (2000a) shows that the real wealth stock backed out using quality-adjusted real investment and geometric, quality-adjusted economic depreciation rates is identical to the productive capital stock. This reflects the fact that the quality-adjusted economic depreciation rate in the simple model equals the rate of physical decay. But Whelan (2000a) notes that the simple model does not allow for the technological obsolescence we observe in the real world: firms sometimes retire productive capital when the marginal product falls below some fixed 'IT support cost'. (Whelan quotes research in the United States by the Gartner Group (1999) that for every $1 spent on computers in 1998, there was another $2.4 spent on wages of IT workers and consultants.) Whelan shows that allowing for such technological obsolescence in the vintage capital model leads to a breakdown of the equivalence between real wealth measures of the capital stock and the productive capital stock. In particular, the economic depreciation rate now exceeds the physical decay rate that should be used in derivation of the productive capital stock. The depreciation rates that we use in our study are economic depreciation rates based on studies underlying the US NIPA measures of the real wealth stock. So on Whelan's arguments they may be too high for growth accounting purposes.

[20] Note that ICT and non-ICT investment includes government investment in these assets respectively. And our measure of the consumption-output ratio includes government consumption.

Table 1: Calibration of baseline model

	Low	High		Low	High		
γ_e	1.150	1.189	α_d	0.264	0.264	τ_l	0.427
γ_z	1.013	1.012	α_e	0.028	0.028	h	0.260
			τ_k	0.298	0.299	β	0.972
δ_d	0.056	0.056				x^d/y	0.136
δ_e	0.241	0.239				x^e/y	0.018

Note: 'Low'/'high' refers to the case where productivity growth in software production is assumed to be low/high.

Finally to determine the remaining parameters, β, α_d, α_e, and τ_k, we estimate the average after-tax real rate of return on capital. We assume that this equals 5.3%, as in Bakhshi, Haldane and Hatch (1999). This is computed using estimates of the 'effective' marginal tax rate on savings in the United Kingdom (which is based on estimates for the average marginal income tax rate on capital income following King and Fullerton (1984) and estimates of the effective tax rate on capital gains). This ties down the ratio β/g. This obviously ties down β for a given estimate of g, but also the three remaining parameters as the solution to the two steady-state Euler equations, (**24**), and the restriction that $(1 - \alpha_d - \alpha_e)$ is equal to labour share of income. The resulting value for α_d is 0.264 and for α_e is 0.028 to a third decimal point for both the 'low software' and 'high software' cases, while the resulting capital tax rate, τ_k, is close to 30% in both cases.

Table 1 summarises the baseline calibration.

3.2 Accounting for growth

We use our 1976-98 sample period for which we have a complete dataset to estimate γ_e. A time series for the ICT capital stock is derived by solving (**4**) recursively, using our time series for Q_t^e, X_t^e and δ_i and an initial value for the capital stock taken from Oulton (2000). A time series for the non-ICT capital stock is derived in a similar fashion. Given our estimates of the capital stocks and α_d and α_e, and the hours worked data, we use the production function in (**13**) to back out a series for Z_t. The annual percentage change in Z_t gives us our estimate of γ_z. With our estimates of γ_z, γ_e and α_e, we can use equation (**20**) to decompose long-run growth into contributions from ICT-specific technological progress and from neutral

Figure 3: ICT-specific and neutral technological progress

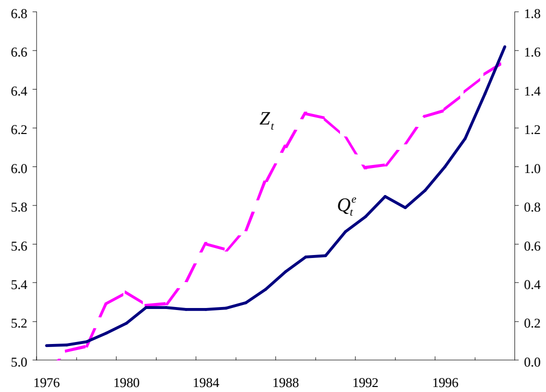

technological progress. The derived series for Z_t is illustrated in Figure 3. Two points can be noted. First, Z_t fell at the beginning of the 1990s, having peaked in 1988 and only recovered to this level in 1997. Second, this period of weak sector-neutral technological progress coincides with the period when investment-specific technological progress for ICT takes off. So when changes in investment-specific technological progress are allowed for, the weakness of TFP in the early 1990s becomes even more pronounced. Of course, movements in Z and Q^e will reflect cyclical as well as trend movements in technological progress.

Table 2 summarises the results from this exercise. ICT-specific technological progress contributes 0.56-0.69 percentage points to labour productivity growth along the balanced growth path, while TFP contributes the remaining 1.75-1.78 percentage points. Note that the implied total productivity growth from this bottom-up exercise is not constrained to equal average labour productivity growth over the sample period: TFP growth is derived by averaging the growth rate of the implied time series Z_t, not by imposing the implied total.

These contributions from ICT-specific technological progress appear very large (around 20%-30% of total labour productivity growth). They reflect the dual assumptions of very sharply falling relative prices of ICT investment goods and the fact that the ICT capital stock as a percentage of GDP in the United Kingdom appears to have been at near-US levels over our sample period. To the extent that the significant contribution of ICT to long-run growth is predicated on sustained falls in the relative price of ICT goods, this echoes the conclusions of Jorgenson and Stiroh (2000), Oliner and Sichel (2000), Gordon (2000), Bosworth and Triplett

Table 2: Decomposing labour productivity growth

	TFP	ICT-specific	Implied total
'Low software'	1.78%	0.56%	2.36%
'High software'	1.75%	0.69%	2.46%

Note: The implied total does not equal actual average labour productivity growth over the 1976-98 sample period. The error reflects differences between the sample average and the balanced growth path.

(2000), and others. They have argued that sustained high productivity growth rates in the United States will in part depend on continued sharp falls in the relative price of computers. This is an important lesson to come out of our balanced growth accounting exercise too.

4 Dynamic aspects of the baseline model

In the previous section, we have characterised the balanced growth path of the model. In the following, we will analyse fluctuations around this steady-state path, caused by temporary but persistent shocks to technology. This analysis is based on a log-linearised approximation to the economy characterised by (**15**)–(**19**), solved using the techniques described in King, Plosser and Rebelo (1988). Using this approximation, we can describe the dynamics of the variables of interest as percentage deviations from the steady-state path described above. Notice that, as before, we assume a constant population so variations in labour inputs are caused by variations in hours. In an economy with deterministic population growth, these variables should be interpreted in per capita terms. The details of these derivations are omitted here, but a technical appendix setting out the details is available on request.

To analyse the effects of shocks to technological progress, a stochastic process for the exogenous shocks must be specified. For that purpose we write

$$Q_t^e = X_t^e q_t^e, \qquad (\mathbf{28})$$

$$Z_t = X_t^z z_t. \qquad (\mathbf{29})$$

where capital letters indicate the trend component and lower-case letters denote the cyclical component. The baseline case that we have used in the growth accounting exercise assumes a deterministic trend so that:

$$\ln\left(X_t^i\right) = \gamma_0^i + t\ln\left(\gamma^i\right),\ i = z, e \qquad \textbf{(30)}$$

To characterise business cycle fluctuations around this trend, we specify that:

$$\begin{aligned} q_t^d &= \exp\left(a_t^d\right);\ \ a_t^d = \rho_d a_{t-1}^d + \varepsilon_t^d,\ i = d, e \\ z_t &= \exp\left(a_{zt}\right);\ \ a_{zt} = \rho_z a_{t-1}^z + \varepsilon_t^z \end{aligned} \qquad \textbf{(31)}$$

We focus solely on impulse response functions, so the only parameters of interest are the persistence parameters ρ_i, $i = e, z$. We estimate ρ_i by fitting an AR(1) with a constant and a linear trend to the series for (the natural logs of) Q_t^e and Z_t derived previously. Depending on the exact measures (low or high software), this exercise suggests that $\rho_e = 0.7$ and $\rho_z = 0.8$ are reasonable values.

We compare the dynamic response to shocks to z_t and to q_t^e. Given the specification chosen here, the shocks we are considering are temporary increases beyond the deterministic trend in productivity: these shocks are persistent, but in both cases, the productivity variable returns to trend. We will later discuss the implication of non-stationary shocks, ie one-off shocks that permanently raise the level of productivity of the economy (though not the growth rate). There are crucial differences in the dynamic responses to these different shocks, but the economic mechanisms are essentially the same.

The impulse-response functions of the baseline model are illustrated in Figure 4. The x-axis of these charts is time, where each period is one year. Shocks occur in period 1. The y-axis is the percentage deviation from the trend path: in the baseline specification, the variables are trend stationary. Both shocks increase investment in capital of type e by increasing the expected marginal product of this type of capital—but while a shock to z_t increases the marginal product of capital on both types of capital, a shock to q_t^e only raises the marginal product of type e. This difference in productivity of capital leads to an immediate re-allocation of capital from production of d to production of e capital: a shock to q_t^e initially causes substitution from investment in d to investment in type e. But the subsequently high e-capital stock raises the marginal product of capital on d-type capital, leading to a subsequent counterflow in investment of the type d: so the initial substitution effect of a shift in relative prices from investment of type d to e is offset in the following periods by a 'complementarity effect' that shifts resources back towards d capital.

To study the response of the aggregate capital stock, the aggregate capital stock is defined as the weighted sum of the two types of capital, where the weights are the

Figure 4: Impulse responses, baseline model

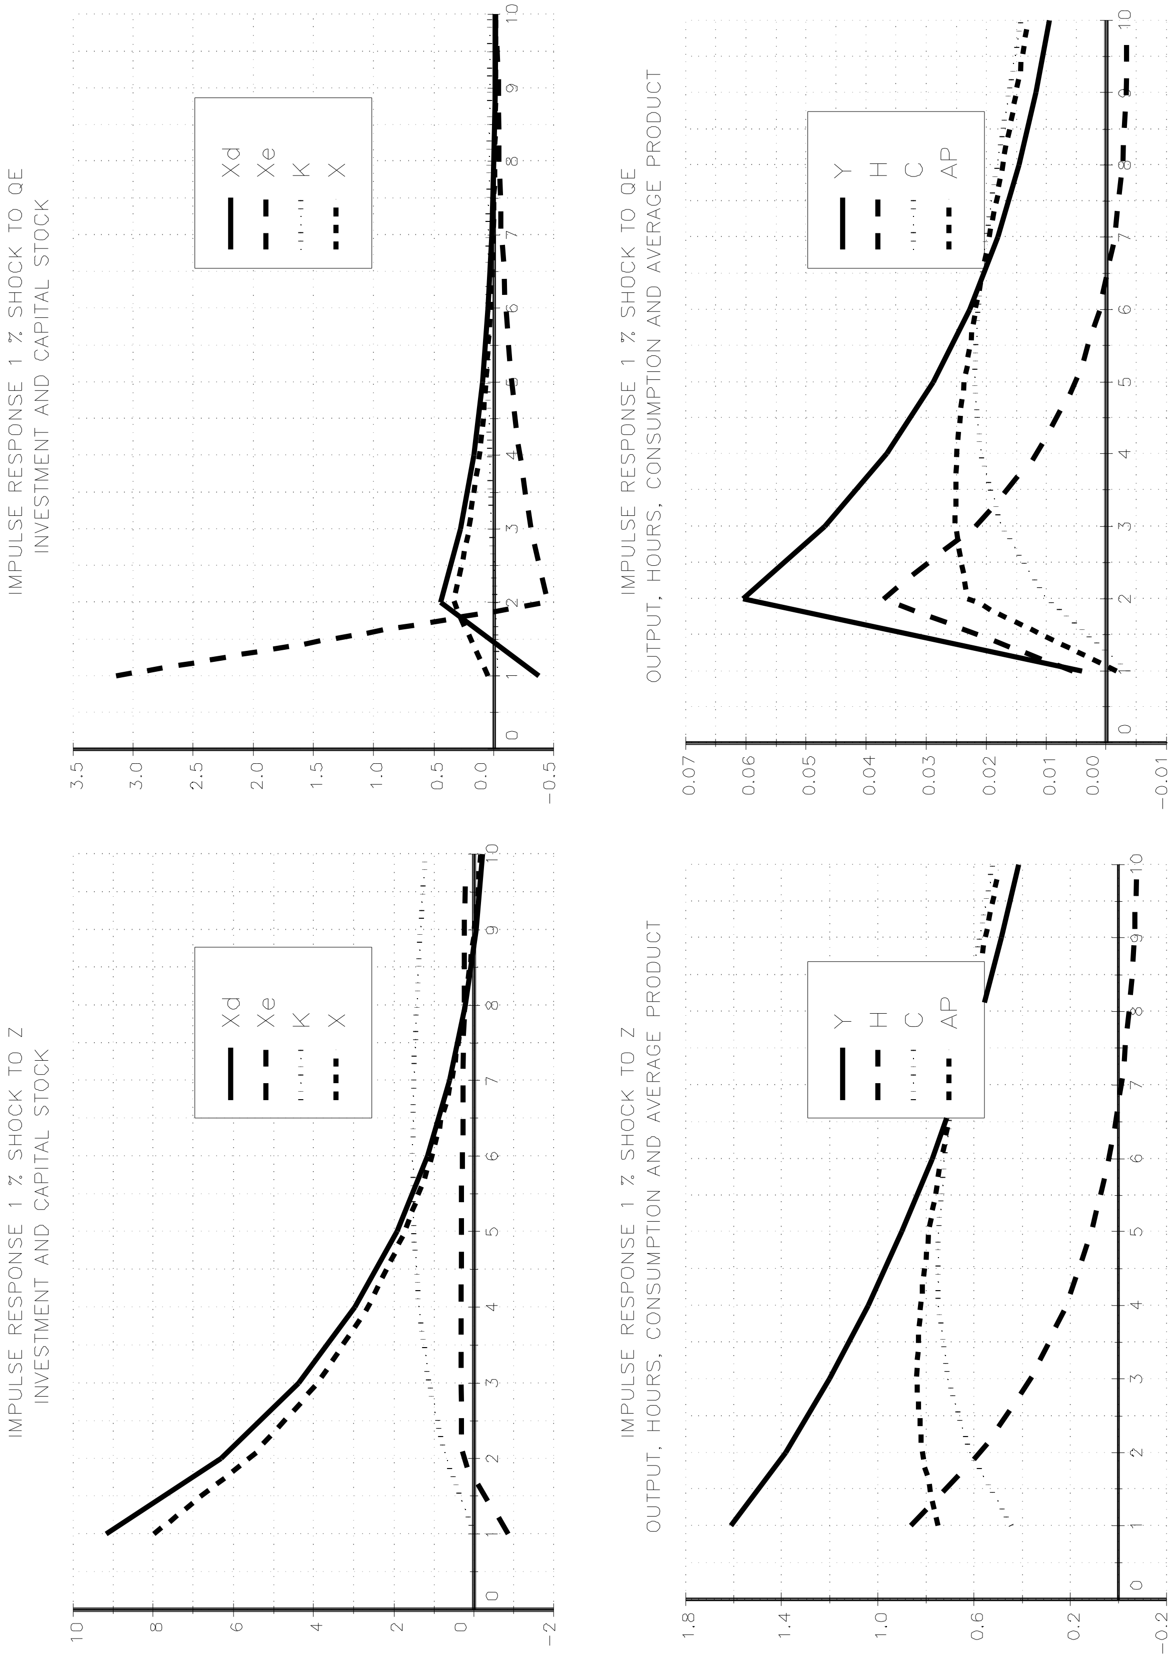

relative prices of capital goods to output:

$$
\begin{aligned}
K_t &= P_t^d K_t^d + P_t^e K_t^e = K_t^d + K_t^e / Q_t^e \\
k_t &= k_t^d + k_t^e / q_t^e
\end{aligned}
\tag{32}
$$

Notice that by assuming that $Q_t^d \equiv 1$, the relative price P_t^d is constant. Note that the aggregate capital stock here is measured in units of final output. So the capital stock K_t grows at the same rate as output, and the output–capital ratio, Y_t/K_t, is stationary. Importantly, K_t is not a state variable: a positive shock to q_t^e lowers the relative price of a component of the capital stock, and hence the replacement value of the entire stock. A shock to z_t, on the other hand, has no such direct effect on the capital stock.

In addition to these differences in the investment response, shocks to q_t^e and to z_t differ in terms of their output implications. A shock to z_t raises output on impact, as more output is produced for given factor inputs. Hours worked also increase as the return to working increases, raising output further; but due to the direct effect of z_t on output, average labour productivity increases on impact. A shock to q_t^e, on the other hand, has no immediate direct effect on output—the effect comes from an increased return to investment, and hence an increase in the capital stock. Output is increased on impact through an increase in hours worked, but this implies a negative rather than a positive effect on average labour productivity. Note how long it takes for labour productivity to settle back to its balanced growth path in both cases. Also, unlike the shock to z_t, the initial effect of a shock to q_t^e on consumption is negative, as resources for extra investment are brought about not only by a decrease in consumption of leisure but also in consumption of goods.

The quantitative effects of the two shocks obviously differ: a shock to q_t^e affects only a small proportion of production and a shock to Z_t is obviously more 'powerful' in the sense that it applies to all production. Yet it is noteworthy that a shock to q_t^e has a stronger effect on output than its share of production would suggest: the peak effect of a 1% shock to ICT-specific technological progress is 0.07%. This suggests that if fluctuations in q_t^e are relatively large, ICT-specific technological progress may account for a large proportion of business cycle fluctuations, despite a relatively small output share. This is in line with Greenwood, Hercowitz and Krusell (2000), who make a similar inference based on technological progress specific to investment in equipment.

5 Extending the baseline model

Above, we have provided a brief characterisation of the baseline model. In this section, we highlight shortcomings of the baseline model as a tool for business cycle analysis. To address these, we modify the model and add features to bring the model more into line with well-known empirical regularities: these do tend to obscure the basic mechanisms discussed in the section above. But the gain is a richer dynamic structure. The features we build in are drawn largely from the existing literature: the main purpose of the exercise is not to provide new theoretical insights, but to analyse the issue at hand, ie sector-specific technological progress, in a model with these features.

One striking feature of the baseline model is that a sector-specific shock causes negative co-movements between sectoral inputs and outputs: a shock to q_t^e leads to an increase in investment of type e but a fall in inputs into production of consumption and d-type investment goods. Similarly, a shock to z_t shifts resources, in the form of hours worked, away from production of consumption goods into production of investment goods (though the net effect on consumption is positive, unlike a shock to investment-specific technological progress). The DGE literature on multi-sector models, reviewed by Greenwood, Hercowitz and Krusell (2000), addresses this issue by including materials (Hornstein and Praschnik (1997)) or intra-sectoral adjustment costs (Huffman and Wynne (1999)). The home production model by Benhabib, Rogerson and Wright (1991) provides a different mechanism to address this issue that is easily implementable in the model considered here: by introducing a home sector to which workers can allocate hours, labour supply to market activities becomes more responsive. In this model, a positive shock to 'market activities', whether investment or sector-neutral technological progress, implies that workers shift hours from the home sector *in addition* to lowering leisure. This issue is analysed in detail in Greenwood, Rogerson and Wright (1995).

One aspect of the baseline model that might appear implausible is the rapid re-allocation of resources from investment in one type of capital to another, or equivalently, the speed with which the capital stock adjusts in response to shocks. The obvious solution in this context is to implement costs to adjusting the capital stock—this, in addition to slowing down the response of the capital stock, affects the response of the price of the capital stock by effectively inserting a wedge between Q_t^i and P_t^i. This is explored in detail in Christiano and Fisher (1995) (who also include habits in consumption). In addition, inclusion of adjustment costs tends to strengthen the propagation mechanism.

The final aspect we look at is variable utilisation rates of capital. Effective capital input then consists of the stock of capital, utilised at a variable rate, with utilisation being costly in the form of increased depreciation. This is important for at least two reasons. Variable utilisation rates imply that effective capital inputs into production can be increased immediately in response to shocks, making output more responsive to shocks and strengthening the propagation mechanism. The implications of this are explored in the literature associated with, amongst others, Burnside and Eichenbaum (1996). Second, a shock to Q_t^i that tends to lower the relative price of ICT capital implies a loss in value for existing capital-holders. And a lower price of the capital stock implies a lower cost, measured in consumption units, of depreciation. This price effect makes it less costly to increase utilisation rates in response to sector-specific shock. This will tend to amplify the output response of a sector-specific shock relative to a neutral shock.

In the following, we provide the details of these extensions to the model. The extensions are implemented in such a way that the steady-state growth path is identical to that of the baseline model. A general property of these extensions is that Q_t^i no longer corresponds directly to the inverse of the deflators. As with the baseline model, we characterise the model by looking at impulse response functions.

5.1 Home production

We introduce home production in the simplest possible way by assuming a home production technology without capital that is linear in hours worked:

$$Y_t^H = Z_t^H h_t^H \tag{33}$$

where Y_t^H is production of home-produced goods, h_t^H is labour input into home production and Z_t^H is labour productivity in the home sector. Home-produced goods are distinct from market goods in that home-produced goods cannot be saved so consumption of these goods necessarily equals production, that is $C_t^H = Y_t^H$. The agent's time constraint is modified to include hours worked at home such that

$$h_t + h_t^H + l_t = 1 \tag{34}$$

As mentioned in the previous section, the model extensions are formulated in such a way that the extended model nests the baseline model. To do so here, we assume the existence of a consumption aggregate $\zeta_t = \zeta\left(C_t, C_t^H\right)$, where ζ is a convex and homogenous aggregator, and write the utility function as $U\left(\zeta_t, L_t\right)$: the baseline specification then simply requires that $\zeta_t = C_t$.

This modification alters the agent's dynamic maximisation problem, adding a first-order condition for hours worked in production at home, and modifying the first-order condition for hours worked in market activities:

$$h_t^H \quad : \quad Z_t^H U_\zeta(\zeta_t, l_t)\, \zeta_{C^H}\left(C_t, C_t^H\right) = U_l(\zeta_t, l_t) \tag{35}$$

$$h_t \quad : \quad U_l(\zeta_t, l_t) = U_\zeta(\zeta_t, l_t)\, \zeta_C\left(C_t, C_t^H\right) W_t (1-\tau_l) \tag{36}$$

The first condition balances the marginal disutility of an extra hour worked with the return to working an additional hour at home, Z_t^H, measured in utility terms, while the second relates the marginal disutility of an extra hour worked with the returns to market activities. These conditions describe how home production alters labour supply; an increase in the real wage now affects labour supply through two channels: it represents not only a decrease in the relative price of market consumption goods relative to leisure but also of market consumption goods relative to the price of home-produced goods. So in this sense, introduction of home production strengthens the substitution effect of an increase in real wages.

To parameterise this extension, only a consumption aggregator is needed. We specify ζ as a CES aggregator

$$\zeta_t = \left\{\left(1-\theta^H\right)(C_t)^e + \theta^H\left(C_t^H\right)^e\right\}^{\frac{1}{e}} \tag{37}$$

which implies that home and market goods are imperfect substitutes with an elasticity of $1/(1-e)$, with θ^H measuring the 'bias' towards home-produced goods. Existence of a steady-state growth path requires that productivity in the home sector Z_t^H grows at the same rate as market output, ie g; this assumption also ensures that the extended model has the same steady-state path as the baseline model. To calibrate the remaining parameters, observe that the first-order condition for allocation of labour implies the following relationship in steady state:

$$\left(\frac{c}{c^H}\right)^{-e}\frac{c}{y}\frac{h}{h^H} = (1-\alpha_d-\alpha_e)(1-\tau_l)\frac{\left(1-\theta^H\right)}{\theta^H} \tag{38}$$

We follow Greenwood, Rogerson and Wright (1995) and set $h^H = 0.25$ and $e = 0$, implying a unit elasticity of substitution between home-produced and market goods. The baseline model's calibration of the remaining parameters and steady-state ratios then imply a value for θ^H.

5.2 Capital

As in the baseline model, we assume that the agent owns the capital stock and rents it to firms on a period-by-period basis. The rental contract specifies an amount of

Figure 5: The depreciation function $g\left(u_t^i\right)$

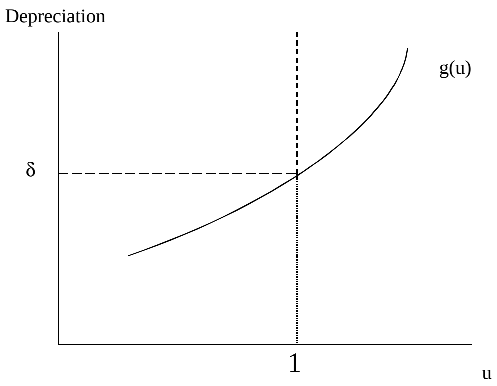

effective capital input, $K_t^{i*} = u_t^i K_t^i$, that the agent will provide to the firm at a fixed price r_t, but the agent determines the composition of the input between utilisation u_t^i and quantities K_t^i. The agent is assumed to determine the capital stock prior to observing the shocks but utilisation after observing the shocks. Increasing utilisation is costly: if the agent decides to increase effective capital supply by increasing utilisation, this results in a higher depreciation rate. Depreciation of capital good i at t is given by

$$\delta_{it} = g_i\left(u_t^i\right),\ i = d, e \tag{39}$$

where g_i is a continuous and convex function $g_i'(.) > 0$ and $g_i''(.) > 0$: increased utilisation of capital increases depreciation at an increasing rate. The properties of the depreciation function are illustrated in Figure 5. In characterising the deviations from steady state, it is the derivative and the elasticity of the derivative that is important, ie how much increases in utilisation translate into increases in depreciation and the elasticity of this response. The baseline case emerges when the elasticity $g_i''(1)u_i/g_i'(1) \to \infty$ (illustrated with the dashed line). In that case, the returns to changing utilisation are becoming infinitely small. This will not affect steady-state utilisation, as this is related to the levels of g, not the derivatives.

In addition to variable utilisation, we introduce a cost of adjusting the capital stock. In particular, we assume the existence of a wedge between investment expenditure measured in units of capital goods, $Q_t^i X_t^i$, and the increase in the capital stock, given by

$$\psi_i\left(\frac{Q_t^i X_t^i}{K_t^i}\right) K_t^i,\ i = d, e$$

The adjustment costs are assumed to be convex by assuming that ψ is concave in its arguments, $\psi_j'(.) > 0$ and $\psi_j''(.) < 0$.

These extensions add an additional first-order condition to the agent's problem, characterising utilisation, and alters the equations that characterise the marginal value of an additional unit of capital and the asset Euler equation. The first-order condition for utilisation dictates that the marginal product of additional utilisation, adjusted for tax and measured in utility terms, equals the marginal cost in the form of increased depreciation:

$$\lambda_t (1-\tau_k) r_t^i K_t^i = \lambda_t^i g_i' \left(u_t^i\right) \frac{K_t^i}{Q_t^i},\ i = d, e \tag{40}$$

where λ_t is the marginal utility of consumption at time t. Here $g_i'(u_t^i)$ is the marginal increase in the rate of depreciation of the capital stock K_t^i. The presence of adjustment costs implies that, out of steady state, there is a wedge between the marginal values of capital stock and the marginal utility:

$$\lambda_t = \lambda_t^d \psi_{dt}' = \lambda_t^e \psi_{et}' \tag{41}$$

Hence the ratio $1/\psi_{it}'$ measures the marginal value in output terms of an additional unit of capital and can be interpreted as a measure of Tobin's marginal q[21] (not to be confused with q_i!): if the derivative of the adjustment cost function is less than 1, this suggests that firms should increase investment (see Figure 6); in the absence of adjustment costs, q is always 1. Notice that the baseline specification requires that the elasticity of ψ' in steady state is 0, while values smaller than 0 indicate more curvature.

Variable utilisation and adjustment costs obviously also modify the Euler asset price equation,

$$\begin{aligned}\frac{\lambda_t^i}{Q_t^i} &= \beta i_t \left(1 - g_i\left(u_t^i\right) + \psi_{it+1} - \psi_{it+1}' \frac{Q_{t+1}^i X_{t+1}^i}{K_{t+1}^i}\right) \frac{\lambda_{t+1}^i}{Q_{t+1}^i} \\ &\quad + \beta i_t \lambda_{t+1}(1-\tau_k) r_{t+1}^i u_{t+1}^i \end{aligned}\tag{42}$$

Finally, the accumulation equation for the individuals' holding of capital is altered to reflect utilisation and adjustment costs:

$$K_{t+1}^i = \left(1 - g_i\left(u_t^i\right)\right) K_t^i - \psi_i \left(\frac{Q_t^i X_t^i}{K_t^i}\right) K_t^i,\ \ i = d, e \tag{43}$$

(21) See Hayashi (1982).

Figure 6: Adjustment cost function

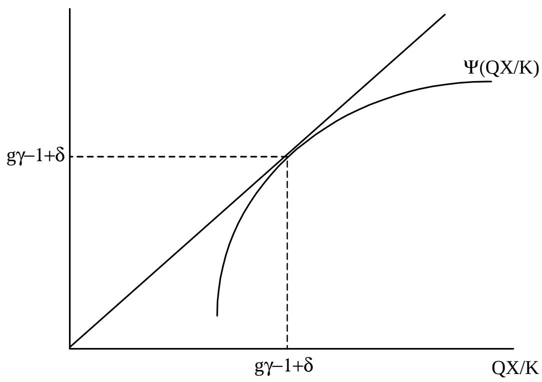

We calibrate the extended model to ensure that the steady-state path is identical to that of the baseline model. To illustrate the restrictions we impose on the adjustment cost function, we return to the sectoral interpretation used previously in Section 2. A capital-producing firm makes output decisions by choosing expenditure on materials M_t, conditional on Q_t^i and K_t^i and given prices P_t^i, to maximise profits $P_t^i \psi_i \left(Q_t^i M_t / K_t^i \right) K_{it} - M_t$. The first-order condition for this problem implies that:

$$P_t^i Q_t^i \psi_i' \left(\frac{Q_t^i X_t^i}{K_t^i} \right) = 1 \tag{44}$$

In the baseline model, the inverse relationship between P_t^i and Q_t^i holds in all equilibria, whether these are on or off the steady-state path. To establish the same relationship on the steady-state path for the extended model, we specify and calibrate the functional form such that $\psi_i' \left(g\gamma_i - 1 + \delta_i \right) = 1$—this implies a steady-state value of one for Tobin's marginal q. Notice that in the extended model the inverse relationship between P_t^i and Q_t^i holds only in steady state.

The accumulation equation function (**43**) imposes an additional restriction on ψ_i: for this accumulation equation to reproduce (**18**) in steady state, we impose that:

$$\psi \left(g\gamma_i - 1 + \delta_i \right) = g\gamma_i - 1 + \delta_i \tag{45}$$

where $g\gamma_i - 1 + \delta_i$ is the investment-capital ratio in the baseline steady state.

In the log-linearised economy, the only additional parameter in ψ_i that needs calibration is the elasticity of ψ_i' in steady state—recall that we have already tied down the level and first derivative in steady state, so we effectively need only to determine the curvature of ψ in the vicinity of steady state. There is no readily available empirical evidence on these two parameters, so we calibrate this parameter by looking at the model's adjustment path when the capital stock is away from steady state. In practice, we set a convergence rate to steady state at 25% a year, implying a half-life of capital stock deviations from steady state of approximately two and a half years. This is roughly equivalent to the values in Basu, Fernald and Shapiro (2000).

For the utilisation function, we impose the restriction that

$$g_i\left(u^i\right) = \delta_i,\ i = d, e \tag{46}$$

ie in steady state, the depreciation rate in the extended model equals that of the baseline model. Notice also that from the first-order condition for utilisation, in steady state:

$$g_i'\left(u^i\right) u^i = \left(1 - \tau_k\right) \alpha_d \frac{y}{k^i} \tag{47}$$

By restricting the depreciation function to be a CES function, (**46**) and (**47**) are sufficient to tie down the necessary parameters.

5.3 Dynamic aspects of the baseline model

As already noted, the extended model is set up in such a way that the steady-state growth path is exactly identical to that of the baseline model—so it only remains to characterise the differences in dynamics around this unchanged steady-state path. As before, we characterise the model by looking at impulse response functions. The impulse responses are illustrated in Figure 7, where the left-hand column of charts are responses to a 1% shock to z_t while the right-hand column are responses to a 1% shock to q_t^e.

As suggested earlier, the presence of adjustment costs implies that there is a wedge between p_t^i and $1/q_t^i$. Prices of capital goods are less responsive to sector-specific shocks because the marginal costs of producing capital goods are sluggish. To illustrate the mechanics of this, recall that marginal costs of producing capital goods are $(q_t^i \psi_t')^{-1}$ and that $\psi'' < 0$. A positive shock to q_t^i directly lowers marginal costs of producing new capital goods, but an increase in production of these goods implies a decrease in ψ_t', offsetting the direct cost effect. This tends to dampen the strong 'substitution effect' seen in the baseline model that leads to a reallocation of resources from production of d to production of e. There is still a 'complementarity effect': a large capital stock of type e raises the marginal product of capital of type d. With a weakened 'substitution effect', this complementarity effect combined with

Figure 7: Impulse responses, extended model

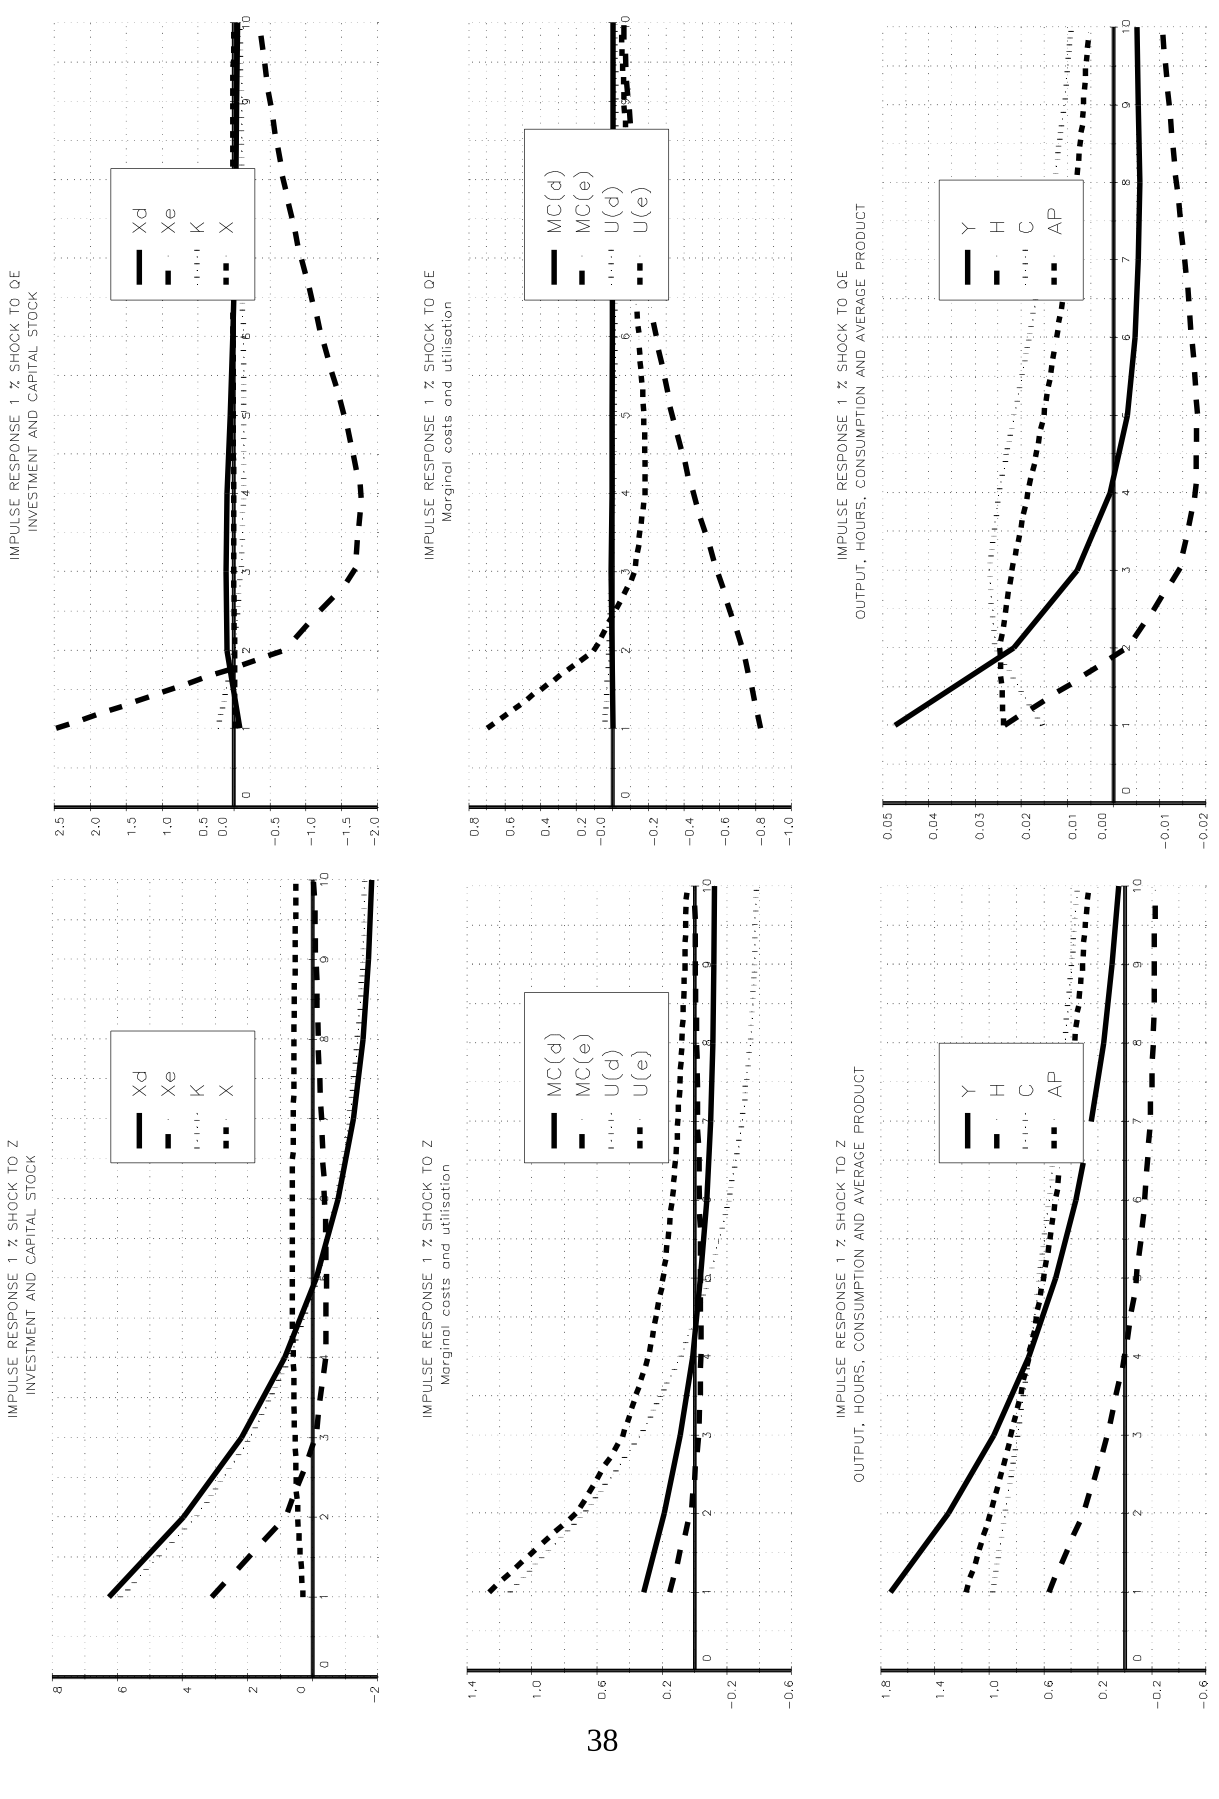

the incentive to smooth investment provided by the adjustment costs implies that investment in capital of type d increases in response to a shock to q_t^e, whereas in the baseline model, the 'substitution effect' dominated (though the effect is quantitatively small). The presence of a home sector reinforces this co-movement: a shock to q_t^e raises the return to market activities relative to home, and this tends to raise production of all market goods. But there are still differences compared with the response of investment in the two types to a sector-neutral z_t shock: the positive co-movement between investment in the two sectors is still much stronger in that case.

The inclusion of adjustment costs also affects the output and consumption responses, primarily by dampening the responses. Importantly, variable utilisation of capital implies that effective capital inputs can be raised immediately in response to shocks—this implies that output can be increased on impact by both increasing hours and utilisation. The return to utilisation increases in response to shocks to both z_t and q_t^e, but a shock to q_t^e has the added effect that the expected future capital loss on existing capital stock, coming from future lower prices, decreases the cost of increased depreciation measured in output terms. The outcome is that in response to a shock to q_t^e, output now increases by more than hours, increasing average labour productivity (unlike in the baseline model).

6 Scenarios for structural change

In the preceding analysis, the maintained assumption has been that a non-stochastic trend is a good description of the economy's steady-state growth path. In this section, we change and relax this assumption in a number of ways. In doing this, we are essentially trying to use the model as tool for 'scenario analysis' of different contemporary examples of structural change. The exercises we consider include temporary and permanent changes to the growth rate of technological progress, as opposed to the temporary changes to the *level* studied in the previous sections. We also look at the implications of changing the technical coefficient α_e on expenditure shares and the aggregate depreciation rate. Our scenario analysis is done in the context of our baseline model, not the extended model of the previous section.

6.1 Permanent shocks to technology

The extensions to the model discussed in Section 5 alter the dynamics around the deterministic path, but maintain the stationary trend assumption. An obvious question is what difference a change in the stochastic properties of the shock would make: Rotemberg and Woodford (1996), for instance, argue in favour of permanent rather than temporary shocks to technology. In the UK context, Ravn (1997) shows that the distinction is important when explaining UK data, but argues that assuming non-stationary shocks *alone* is insufficient when explaining business cycle facts. While being somewhat agnostic on this issue, we want to consider the implications of permanent shocks to technological progress in our model. Arguably, such a shock is a better characterisation of the views of some proponents of the 'New Economy' hypothesis: they argue that the US economy may have experienced an increase in medium-term productivity growth, but that it is still too early to conclude anything about long-run growth. And on this view, a temporary shock to the growth rate of productivity might be a more pertinent simulation for policy-makers in the United Kingdom who wish to embed a 'New Economy' shock into their macroeconomic forecasts. We do this by modifying (**30**) so that:

$$\ln\left(X_t^i\right) = \ln\left(\gamma^i\right) + \ln\left(X_{t-1}^i\right) + \varepsilon_t^i,\ i = e, z \qquad (\mathbf{48})$$

We set the drift parameter such that the average growth rates of the model with non-stationary shocks are identical to those of the baseline model, so that in the absence of shocks the two economies would follow the same growth path. Furthermore, we assume that there are no temporary shocks.[(22)] So a shock in this new economy shifts the level of productivity permanently, whereas in the baseline model, a shock has only temporary (though persistent) effect. We analyse this issue using the baseline specification of the model.

This change in the stochastic properties of the shocks changes the normalisation of variables. We replace the terms g^t, g_d^t and g_e^t that characterise the non-stochastic steady-state path with three stochastic terms:

$$N_t = Z_t^{\frac{1}{\alpha_l}} \left(Q_t^e\right)^{\frac{\alpha_e}{\alpha_l}};\ N_t^d = N_t;\ N_t^e = Q_t^e N_t \qquad (\mathbf{49})$$

where $\alpha_l = 1 - \alpha_e - \alpha_d$. The variables are now normalised as follows:[(23)]

$$\begin{aligned} y_t &= Y_t/N_t;\ c_t = C_t/N_t;\ x_t^e = X_t^e/N_t;\ x_t^d = X_t^d/N_t;\ k_t^d = K_t^d/N_{t-1}^d;\ k_t^e = K_t^e/N_{t-1}^e; \\ \gamma_t^d &= 1;\ \gamma_t^e = Q_t^e/Q_{t-1}^e;\ \gamma_t^z = Z_t/Z_{t-1};\ \tilde{\lambda}_t^e = \lambda_t^e N_t;\ \tilde{\lambda}_t^d = \lambda_t^d N_t \qquad (\mathbf{50}) \end{aligned}$$

---

(22) By assuming that there are only permanent shocks, we avoid a potential signal extraction problem: if there were both permanent and temporary shocks, then the agents would have to separately identify the shocks.

(23) Notice that the growth rates γ_t^d, γ_t^e and γ_t^z are now stochastic. Also, to accommodate a stochastic growth rate, we have changed the timing convention on the capital stock normalisation.

The stationary system is then characterised by the following sets of equations

$$\begin{aligned}
\frac{h_t}{1-h_t} &= \frac{\theta}{1-\theta}(1-\tau_l)(1-\alpha_e-\alpha_d)\frac{y_t}{c_t} \\
\tilde{\lambda}_t^e &= \tilde{\lambda}_t^d = \frac{\theta}{\tilde{C}_t} \\
c_t + x_t^e + x_t^d &= \left(\frac{1}{\gamma_t^N \gamma_t^d}\right)^{\alpha_d} \left(\frac{1}{\gamma_t^N \gamma_t^e}\right)^{\alpha_e} (k_t^e)^{\alpha_e} \left(k_t^d\right)^{\alpha_d} h_t^{1-\alpha_e-\alpha_d} \\
\tilde{\lambda}_t^i &= \beta \tilde{\lambda}_{t+1}^i (1-\tau_k)\alpha_i \frac{y_{t+1}}{k_{t+1}^i} + \beta \tilde{\lambda}_{t+1}^i (1-\delta_i) \frac{1}{\gamma_{t+1}^N \gamma_{t+1}^i} \\
k_{t+1}^i &= (1-\delta_i) k_t^i \frac{1}{\gamma_t^N \gamma_t^i} + x_t^i, \quad i = d, e
\end{aligned} \tag{51}$$

where

$$\gamma_t^N = (\gamma_t^z)^{\frac{1}{\alpha_l}} (\gamma_t^e)^{\frac{\alpha_e}{\alpha_l}} \tag{52}$$

As mentioned, the shocks now permanently change the steady-state level of output. From (**49**), a 1% permanent increase in the level of Z_t permanently raises the steady-state path of output by $1/\alpha_l\%$, whereas a similar shock to Q_t^e shifts the steady-state path by α_e/α_l. The dynamics of the adjustment paths are illustrated in Figure 8—the economic mechanisms discussed at length in Section 4 are the same, but the dynamics differ. Unlike a neutral shock, a permanent increase in Q_t^e initially *decreases* output as hours worked fall. We ascribe this to the income effect dominating the substitution effect: having observed a permanent shock to technology, whether sector-specific or neutral, agents will know that long-run income levels have increased. This tends to lower labour supply. In the case of a sector-neutral shock, there is a strong offsetting substitution effect from an immediate increase in wages (or equivalently, an increase in the cost of leisure). With a sector-specific shock, there is no such effect in the first period because increases in Q_t^e do not affect output on impact. Productivity and hence wages increase only in subsequent periods, which then increases labour supply.

The income/substitution effects also distinguish the investment/consumption responses in the two cases. Here the counterbalancing is between a falling price of investment goods or an increasing return to investment on the one hand (substitution effect), and a permanent increase in income that tends to increase consumption at the expense of investment on the other (income effect). With a neutral shock, the return to investment increases for both types of capital good. This dominates the income effect, so aggregate investment overshoots its long-run levels and the consumption-output ratio decreases. A shock that is specific to production of investment goods of type e only raises the return to investment in capital of this type: it shifts resources from production of type d goods, but aggregate investment undershoots its long-run level as the income effect dominates the substitution effect and the consumption-output ratio increases.

6.2 ICT investment expenditure share

Even with permanent shocks to productivity growth, the balanced growth path is characterised by constant expenditure shares: production becomes increasingly ICT intensive but the price of this capital good is falling, leaving the investment expenditure share constant. Arguably, one feature of the recent US experience is a sharp increase in the ICT investment expenditure share—certainly in the United Kingdom, the investment expenditure share has been rising sharply over the period, with the ICT share increasing from 0.7% in 1976 to 3.6% in 1998.

Accounting for this phenomenon poses a challenge to the model we are using. To some extent, the model can account for this as a *temporary* phenomenon: in the baseline model, a fall in the price of ICT-capital goods leads to a large temporary increase in investment that exceeds the fall in prices, so the ICT-investment expenditure share temporarily increases. But with very rapid adjustment of factor inputs, the steady-state expenditure share is quickly restored. The extensions of the baseline model dampen and slow down this adjustment, implying a smaller but more persistent response of the investment expenditure share. The baseline model cannot account for this as a *permanent* phenomenon: along the steady-state growth path where growth is balanced the expenditure shares are constant. To analyse *permanent* changes to the investment expenditure share, we need to consider changes in structural parameters. In the following, we perform some comparative statics exercises and characterise how changes in some structural parameters change the balanced growth path, holding all other parameters at their steady-state values.

The obvious first candidate is to change the growth rates of sector-specific technological progress, ie to consider changes to γ^e in (**48**), similar to the exercise in Pakko (2000). From (**52**), an increase in γ^e also increases the aggregate growth rate γ^N. Such a change has two offsetting effects on the investment expenditure share. To see this, consider the steady-state version of the capital accumulation and the Euler equations from (**51**):[24]

$$\begin{aligned} (1-\tau_k)\,\alpha_e \frac{y}{k^e} + \frac{(1-\delta_e)}{\gamma^N \gamma^e} &= \frac{1}{\beta} \\ 1 - \frac{(1-\delta_e)}{\gamma^N \gamma^e} &= \frac{x^e}{k^e} \end{aligned} \qquad (\mathbf{53})$$

[24] Recall that the capital stock of type i is normalised on $N_{t-1}Q^i_{t-1}$.

Figure 8: Impulse responses, permanent shocks

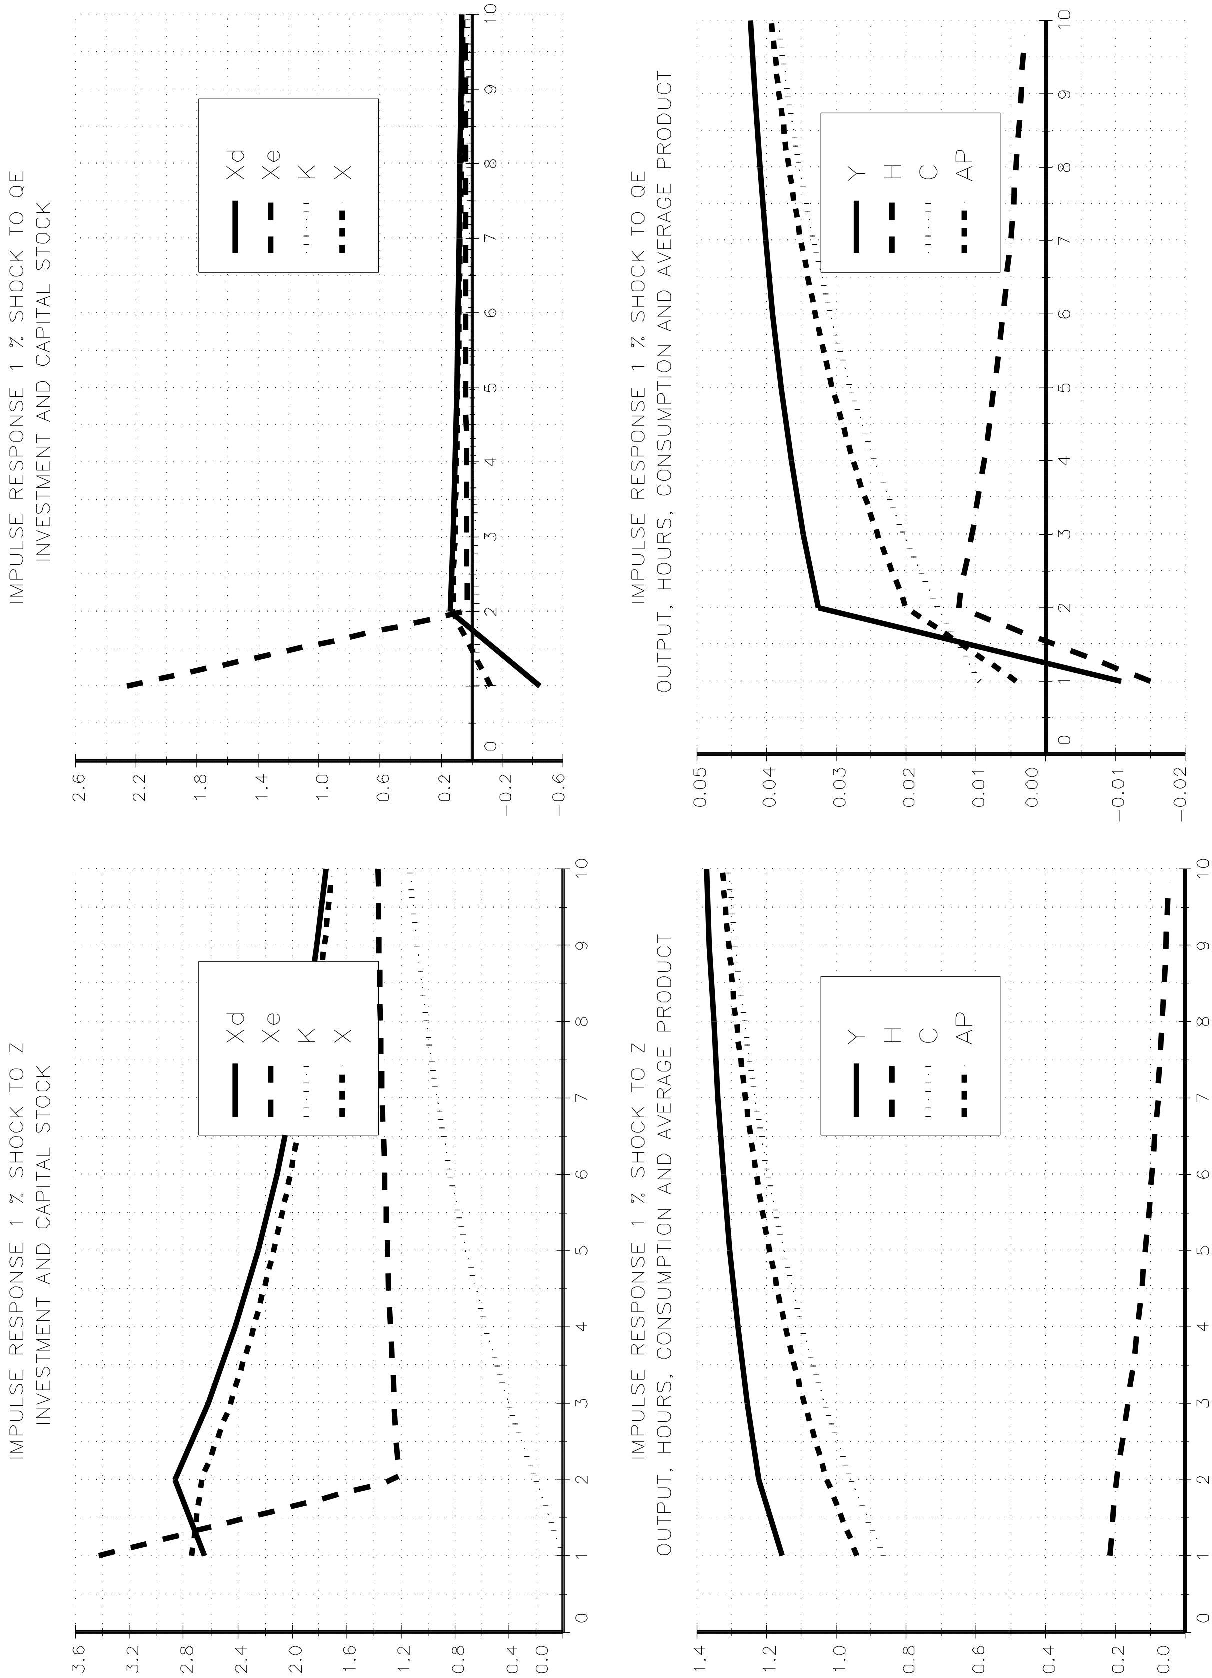

An increase in the growth rate γ^e leads to an increase in the y/k_e ratio through a negative 'capitalisation effect': the return to investing in one unit of capital is the after-tax marginal product of capital, plus the value of the capital stock next period, $(1-\delta_e)/\gamma^N\gamma^e$. An increase in γ^e lowers the value of the capital, because the intertemporal price of capital good e is falling faster. For a given discount factor, this will require an increase in the return to capital, ie an increase in the y/k^e ratio to increase the marginal product of capital. On the other hand, there is an 'accumulation effect': an increase in the growth rate will require an increase in the investment-capital ratio, x^e/k^e, to maintain balanced growth. In combination, the ratio x^e/y is given by:

$$\frac{x^e}{y} = \frac{x^e}{k^e}\frac{k^e}{y} = \frac{\gamma^N\gamma^e - (1-\delta_e)}{\gamma^N\gamma^e - \beta(1-\delta_e)}\beta(1-\tau_k)\alpha_e \tag{54}$$

It is straightforward to establish that, provided $\beta < 1$, x^e/y is increasing in γ^e, so in other words, the accumulation effect dominates. Figure 9 depicts this relationship: notice that even substantial changes in γ^e (the x-axis) lead to fairly small changes in expenditure share (the y-axis). Hence, an increase in γ^e to match the increased ICT investment expenditure ratios would require the growth rate γ^e to increase substantially, implying in turn a substantial increase in the steady-state growth rate.(25)

Figure 9 also shows the effect of varying γ^e on the aggregate depreciation rate: an increase in the growth rate of ICT-specific technological progress would imply a *decrease* in the aggregate depreciation rate. So despite the fact that an increase in γ^e leads to higher growth in intensity of a capital good with a relatively higher depreciation rate, the aggregate depreciation rate falls. This, essentially, is caused by the capitalisation effect. To see this, we define the aggregate depreciation rate as

$$\delta_t = \omega_{dt}\delta_d + \omega_{et}\delta_e \tag{55}$$

where the weights ω_{it} are constant price shares in aggregate capital stock.(26) This

(25)Notice that changes in x^i/y that lead to changes in c/y will affect hours worked h. First, $\frac{c}{y} = 1 - \sum_i \frac{x^i}{y}$; using this, $h = \left(\frac{1-\tau_t}{\alpha_l}\frac{\theta}{1-\theta}\frac{c}{y}\right) / \left(1 + \frac{1-\tau_t}{\alpha_l}\frac{\theta}{1-\theta}\frac{c}{y}\right)$.

(26)Note that we use constant price capital stocks here to weight together ICT and non-ICT depreciation rates, compared with current price weights when we calibrate the average ICT and non-ICT depreciation rates. This is because we will want to compare the resulting average depreciation rate with those implied from official non-chain weighted capital stocks. In contrast, our own capital stock data that we use to calibrate the model are chain-weighted. Tevlin and Whelan (2000) show that severe biases can arise when comparing depreciation rates implied from chain-weighted and non-chain-weighted data.

Figure 9: Effects of variations in g_e

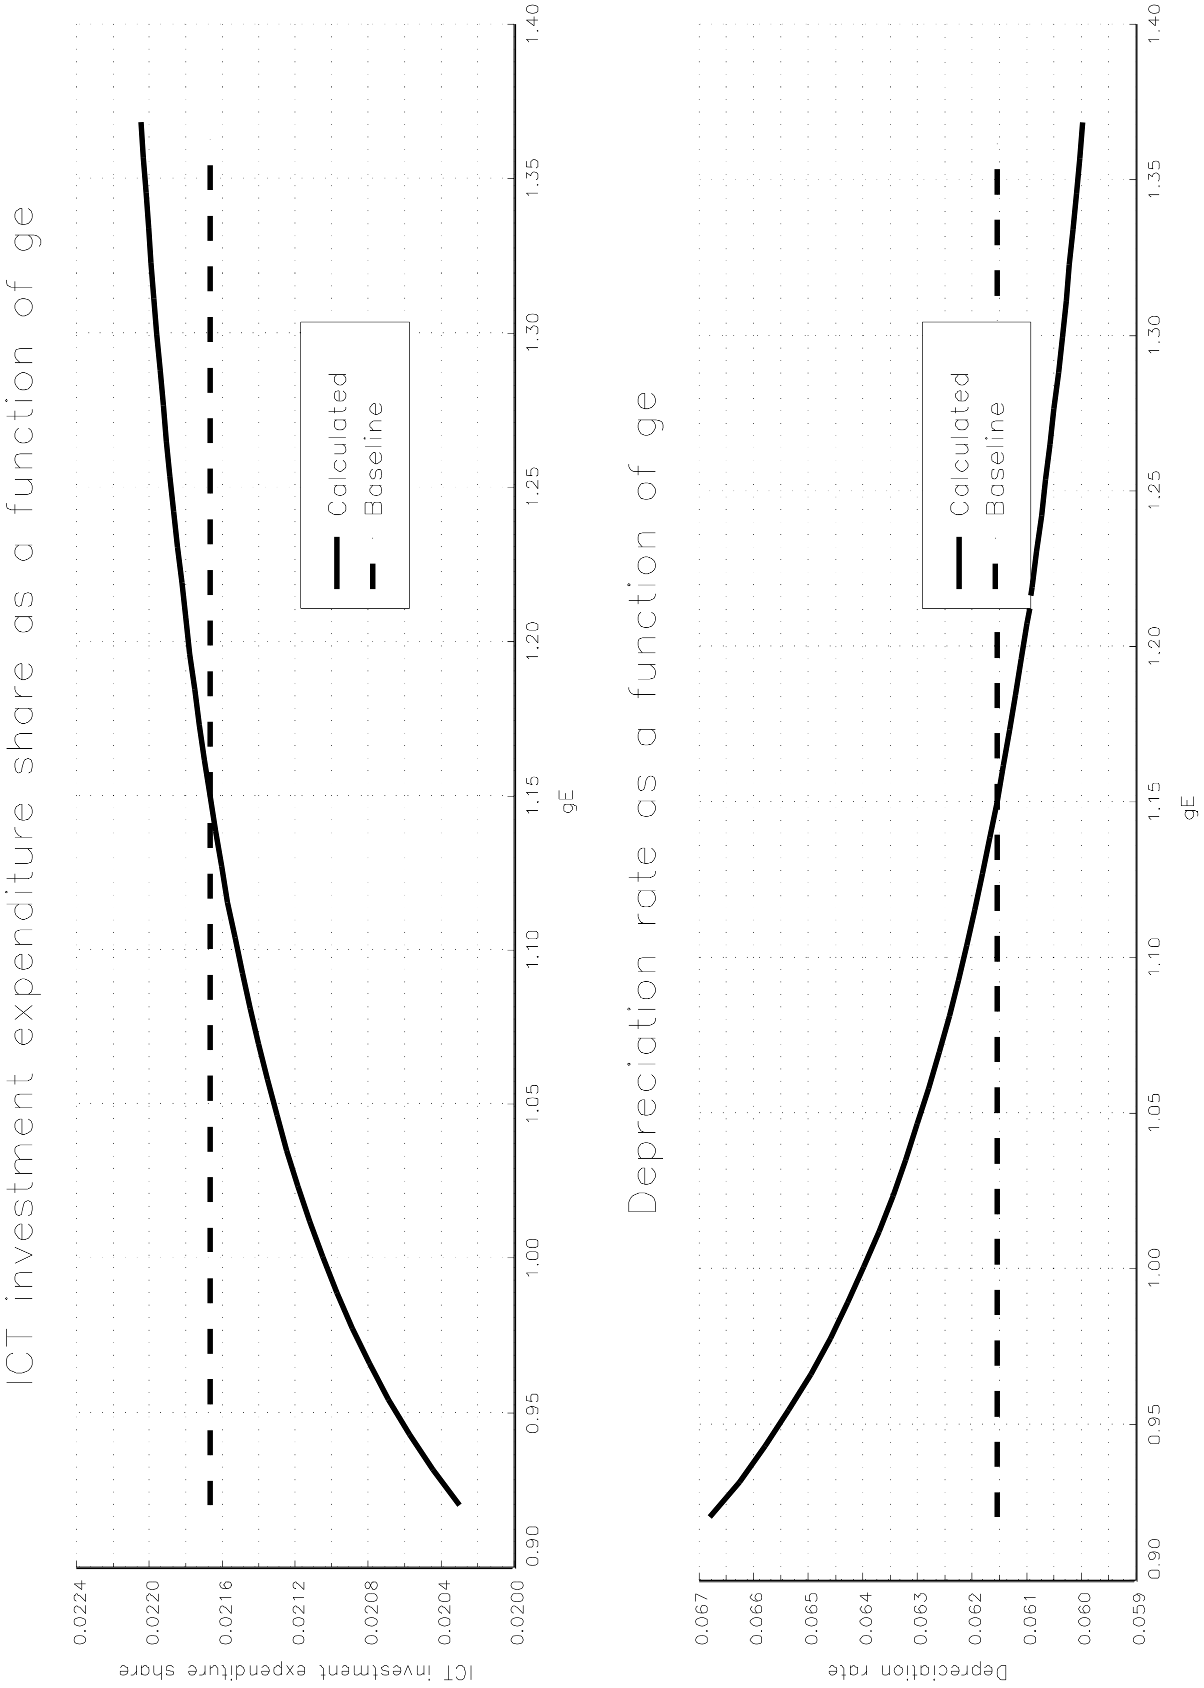

we can write as

$$\begin{aligned}\delta_t &= \frac{K_t^d}{K_t}\delta_d + \frac{K_t^e/Q_t^e}{K_t}\delta_e \\ &= \frac{k_t^d/y_t}{k_t/y_t}\frac{1}{\gamma_t^N}\delta_d + \frac{k_t^e/y_t}{k_t/y_t}\frac{1}{\gamma_t^e\gamma_t^N}\delta_e \end{aligned} \tag{56}$$

where the aggregate capital stock and capital–output ratios are defined as:[27]

$$\begin{aligned}K_t &= P_t^d K_t^d + P_t^e K_t^e, \\ \frac{k_t}{y_t} &= \frac{k_t^d}{y_t}\frac{1}{\gamma_t^N} + \frac{k_t^e}{y_t}\frac{1}{\gamma_t^e\gamma_t^N}\end{aligned} \tag{57}$$

From (**56**), an increase in γ^e lowers the weight on δ^e, which would tend to decrease the aggregate depreciation rate. But a change in γ_t^e also affects the weight on k_t^d, as γ_t^e has a direct effect on γ_t^N, so the net effect will depend on the exact calibration of the model. Figure 9 shows that under the current calibration, the weight on δ^e falls more than that the weight on δ^d, leading to a decrease in the aggregate depreciation rate. This means that an increase of γ^e is inconsistent with the empirical evidence on depreciation rates. Official investment and capital stock data at the plant and machinery level are available for both the United Kingdom and United States.[28] These can be used to back out implied depreciation rates. Figure 10 shows that the implied depreciation rates for both the United Kingdom and United States (indexed to 1990=100) have increased since 1990.

In summary, an increase in γ^e can increase the investment expenditure share, but accounting for the observed increase would require a substantial increase in γ^e. In addition, such an increase lowers the aggregate depreciation rate, which is at odds with the empirical evidence.

A direct change in the technical parameter α_e also increases the investment expenditure share and, contrary to the previous experiment, *increases* the aggregate depreciation rate. In the experiment we consider here, we increase α_e but hold $\alpha_e + \alpha_d$ constant—ie an increase in α_e is offset by a decrease in α_d. By calculating the derivative of γ^N with respect to α_e under this assumption, it is straightforward to establish that the steady-state growth rate γ^N increases with an increase in α_e.

(27) Notice that while K_t^d and K_t^e are state variables, the aggregate capital stock K_t is not: the fact that K_t is measured in units of final goods means that K_t can change instantaneously in response to shocks. For this reason, K_t is normalised on N_t rather than N_{t-1}.

(28) We are grateful to Stacey Tevlin for providing us with the US data. The implied rates are calculated using a fixed-weight measure of the capital stock. These data include computers and communications equipment, but exclude software.

Figure 10: Implied depreciation rates for plant and machinery

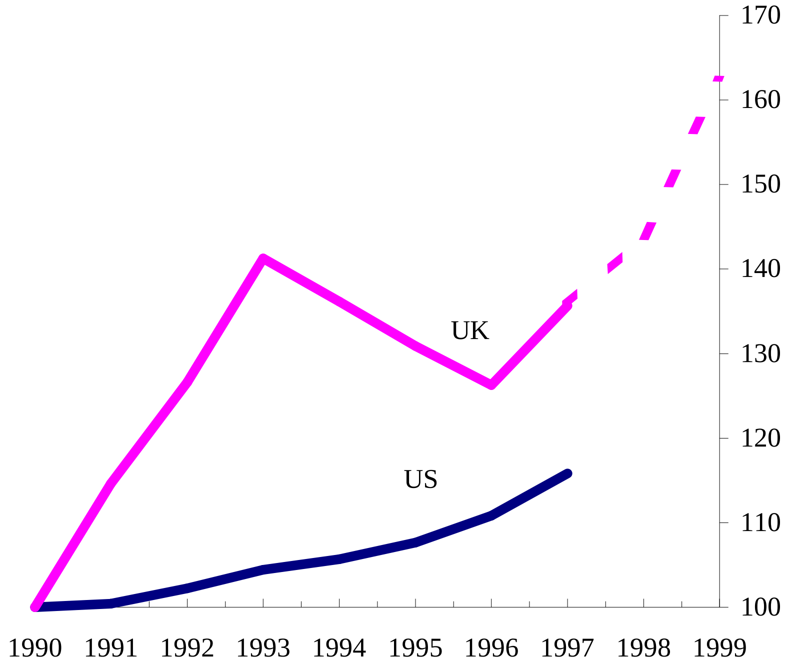

So while there are still capitalisation and accumulation effects, stemming from increases in growth rates, the capitalisation effect is now offset by a direct increase in the return to investment. Figure 11 draws out the change in steady-state investment expenditure ratios and depreciation rates, as a function of a change in α_e, holding $\alpha_d + \alpha_e$ constant. To increase the ICT investment share of output to match the last observation in our dataset, α_e should be increased to 0.054, from a benchmark value of 0.031. This implies an increase in the depreciation rate from the steady-state value of 6.45% to 6.9%, or an approximately 7% increase. This increases the growth rate of output to 2.6%.[(29)]

7 Conclusion

In this paper we have decomposed labour productivity growth along the balanced growth path of a model of the UK economy into investment-specific and sector-neutral technological progress. Using US hedonic deflators for ICT investment goods, we find that ICT investment-specific technological progress

[(29)] One could obviously consider a CES rather than a Cobb-Douglas aggregator. An increase in the ICT share of the aggregate capital stock would then require that the elasticity of substitution between the two types of capital was *greater* than one.

Figure 11: Effects of variations in a_e

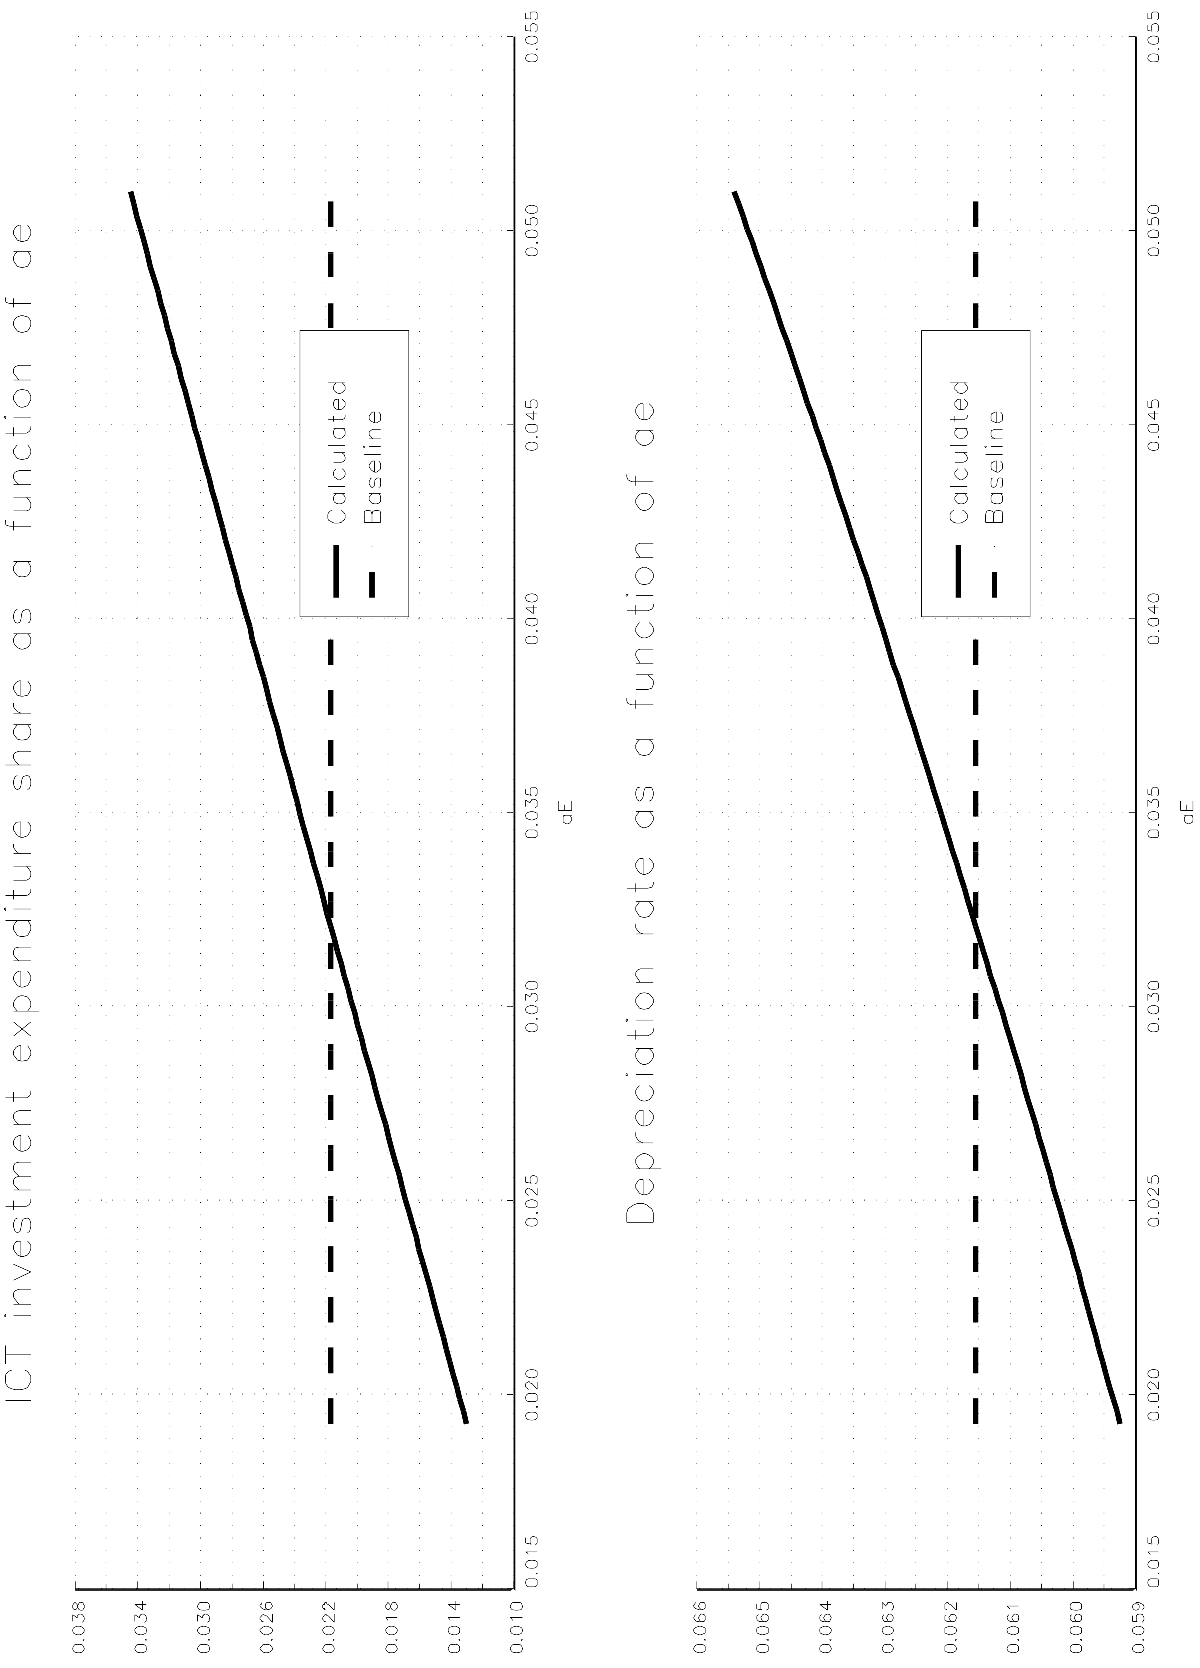

makes a significant contribution to productivity growth along the balanced growth path, explaining as much as 20%-30% of labour productivity growth. One obvious conclusion is that sustained improvements in labour productivity growth from this source will rely on continued sharp declines in the relative price of ICT goods.

We have drawn out the different implications of shocks to investment-specific technological progress on the one hand, and sector-neutral technological progress on the other. Such differences are important for policy-makers who wish to embody future 'New Economy' productivity shocks into their macroeconomic forecasts. In addition to this dynamic analysis, we have also performed some comparative statics exercises, characterising how the balanced growth path is affected by changes in underlying parameters. We have not, in this paper, considered the exact dynamics of how the economy might move from one balanced growth path to another, although the model can obviously be used for such an exercise.

References

Bakhshi, H, Haldane, A and Hatch, N (1999), 'Some costs and benefits of price stability in the UK', in Feldstein, M (ed) *The Costs and Benefits of Price Stability.*

Basu, S, Fernald, J G and Shapiro, M (2000), 'Productivity growth in the 1990s: technology, utilization, or adjustment?', *mimeo*.

Benhabib, J, Rogerson, R and Wright, R (1991), 'Homework in macroeconomics: household production and aggregate fluctuations', *Journal of Political Economy*, Vol.99, pages 1,166–87.

Bosworth, B and Triplett, J (2000), 'What's new about the new Economy? ICT, economic growth and productivity', paper presented at 'Technical Progress, Economic Growth and the New Economy', National Institute of Economic and Social Research, United Kingdom.

Broadbent, B and Walton, D (2000), 'How 'New' is the UK economy?', *Goldman Sachs Global Economics Papers, no.43*.

Brookes, M, and Wahhaj, Z (2000), 'The shocking economic effect of B2B', *Goldman Sachs Global Economics Paper, no.37*.

Brynjolfsson, E and Kemerer, C (1996), 'Network externalities in microcomputer software: an econometric analysis of the spreadsheet market', *Management Science*, Vol.42, pages 1,627–47.

Burnside, C and Eichenbaum, M (1996), 'Factor-hoarding and the propagation of business-cycle shocks', *American Economic Review*, Vol.86(5), pages 1154–74.

Christiano, L J and Fisher, J (1995), 'Tobin's Q and asset returns: implications for business cycle analysis', *Discussion Paper 5292*, NBER.

Fraumeni, B (1997), 'The measurement of depreciation in the US national income and product accounts', *Survey of Current Business*, pages 7–23.

Gartner Group (1999), '1998 IT spending and staffing survey results', *Gartner Group*.

Gordon, R J (1990), *The Measurement of Durable Goods Prices*, University of Chicago Press.

Gordon, R J (1999), 'Has the 'New Economy' rendered the productivity slowdown obsolete?', *mimeo*.

Gordon, R J (2000), 'Does the 'New Economy' measure up to the great inventions of the past?', *mimeo*.

Gort, M, Greenwood, J and Rupert, P (1999), 'Measuring the rate of technical progress in structures', *Review of Economic Dynamics*, Vol.2, pages 207–30.

Greenwood, J, Hercowitz, Z and Huffman, G W (1988), 'Investment, capacity utilization, and the real business cycle', *American Economic Review*, Vol.78(3), pages 402–17.

Greenwood, J, Hercowitz, Z and Krusell, P (1997), 'Long-run implications of investment-specific technological change', *American Economic Review*, vol.87, pages 342–62.

Greenwood, J, Hercowitz, Z and Krusell, P (2000), 'The role of investment-specific technological change in the business cycle', *European Economic Review*, Vol.44, pages 91–115.

Greenwood, J, Rogerson, R and Wright, R (1995), 'Household production in real business cycle theory', in Cooley, T F (ed) *Frontiers of Business Cycle Theory*, Princeton University Press.

Gust, C and Marquez, J (2000), 'Productivity developments abroad', *Federal Reserve Bulletin*, pages 665–81.

Hayashi, F (1982), 'Tobin's marginal and average q: a neoclassical interpretation', *Econometrica*, pages 213–24.

Hercowitz, Z (1998), 'The 'embodiment' controversy: a review essay', *Journal of Monetary Economics*, Vol.41, pages 217–24.

Hornstein, A and Krusell, P (1996), 'Can technology improvements cause productivity slowdowns?', *NBER Macroeconomics Annual*, Vol.11, pages 209–59.

Hornstein, A and Praschnik, P (1997), 'Intermediate inputs and sectoral comovements in the business cycle', *Journal of Monetary Economics*, Vol.40, pages 573–95.

Huffman, G W and Wynne, M A (1999), 'The role of intratemporal adjustment costs in a multisector economy', *Journal of Monetary Economics*, Vol.43, pages 317–50.

Hulten, C R (1978), 'Growth accounting with intermediate inputs', *Review of Economic Studies*, Vol.45, pages 511–18.

Hulten, C R (1979), 'On the 'importance' of productivity change', *American Economic Review*, Vol.69, pages 126–36.

Hulten, C R (1992), 'Growth accounting when technical change is embodied in capital', *American Economic Review*, Vol.82, pages 964–80.

Jenkins, S and O'Leary, M (1997), 'Gender differentials in domestic work, market work, and total work time: UK time budget survey evidence for 1974/5 and 1987', *Scottish Journal of Political Economy*, Vol.44, pages 153–64.

Jorgenson, D W (1966), 'The embodiment hypothesis', *Journal of Political Economy*, LXXIV, pages 1–17.

Jorgenson, D W (1995), *Productivity, Postwar US Economic Growth.* MIT Press, Cambridge, MA.

Jorgenson, D W and Stiroh, K J (2000), 'Raising the speed limit: US economic growth in the information age', *mimeo*.

King, M and Fullerton, D (1984), *The Taxation of Income from Capital: A Comparative Study of the US, UK, Sweden and West Germany*, University of Chicago Press.

King, R G, Plosser, C I and Rebelo, S T (1988), 'Production, growth and business cycles. II. New directions', *Journal of Monetary Economics*, Vol.21, pages 309–41.

Kneller, R and Young, G (2000), 'The new British economy: the record so far and the future outlook', paper presented at 'Technical Progress, Economic Growth and the New Economy', National Institute of Economic and Social Research, United Kingdom.

Kydland, F E and Prescott, E C (1982), 'Time to build and aggregate fluctuations', *Econometrica*, Vol.50, pages 1,345–70.

Millard, S P, Scott, A and Sensier, M (1997), 'The labour market over the business cycle: can theory fit the facts?', *Oxford Review of Economic Policy*, Vol.13(3), pages 70–92.

Moulton, B, Parker, R and Seskin, E (1999), 'A preview of the 1999 comprehensive revisions of the national income and product accounts', *Survey of Current Business*, pages 7–20.

Oliner, S D and Sichel, D E (1994), 'Computers and output growth revisited: how big is the puzzle', *Brookings Papers on Economic Activity*, Vol.2, pages 273–334.

Oliner, S D and Sichel, D E (2000), 'The resurgence of growth in the late 1990s: is information technology the story?', *Journal of Economic Perspectives*.

Oulton, N (2000), 'ICT and productivity growth in the UK', Bank of England, *mimeo*.

Pakko, M R (2000), 'What happens when the technology growth trend changes? Transition dynamics and growth accounting', Federal Reserve Bank of St. Louis *Discussion paper.*

Parker, R and Grimm, B (2000), 'Software Prices and Real Output: Recent Developments at the BEA', paper presented at the NBER programme on 'Technological Change and Measurement'.

Ravn, M O (1997), 'Permanent and transistory shocks, and the UK business cycle', *Journal of Applied Econometrics*, Vol.12(1), pages 27–48.

Rotemberg, J J and Woodford, M (1996), 'Real-business-cycle models and the forecastable movements in output, hours, and consumption', *American Economic Review*, Vol.86(1), pages 71–89.

Solow, R M (1960), 'Investment and Technical Progress', in Arrow, K J, Karlin, S and Suppes, P (eds) *Mathematical Methods in the Social Sciences*, Stanford University Press.

Tevlin, S and Whelan, K (2000), 'Explaining the investment boom of the 1990s', *mimeo*.

Whelan, K (2000a), 'Computer, obsolescence, and productivity', Federal Reserve Board Discussion paper.

Whelan, K (2000b), 'A guide to the use of chain-aggregated NIPA data', Federal Reserve Board *discussion paper.*