How well do we understand business cycles and growth? Examining the data with a real business cycle model.*

Harald Uhlig
Humboldt University Berlin
CentER, Tilburg University
and CEPR

First draft: August 2002
This revision: January 20, 2003

Abstract

How well do we understand business cycles and growth? One benchmark answer is provided by the stochastic neoclassical growth theory or real business cycle theory: growth results from the steady progress of productivity, while business cycles result from fluctuations of productivity around this trend. In this paper, we examine how well a rather simple real business cycle model fits the key facts. The fit at business cycle frequencies is quite compelling, given the simplicity of the model. However, the model does not seem to come much closer to the data than a "primitive" model, which just rescales TFP fluctuations. Moreover, the fit becomes less convincing, when no filtering is done. The challenge is to construct better quantitative theories that address these deficiencies.

Keywords: real business cycle, total factor productivity, quantitative macroeconomics

JEL codes: C10, C63, C82, E13, E20, E32

*I am grateful to two discussants, Jörg Breitung and Gerd Ronning, as well as to seminar participants at Ottobeuren 2002 for many useful comments, and to Mathias Trabandt for a careful reading of the paper. All remaining errors are my own. Address: Prof. Harald Uhlig, Humboldt University, Wirtschaftswissenschaftliche Fakultät, Spandauer Str. 1, 10178 Berlin, GERMANY. e-mail: uhlig@wiwi.hu-berlin.de, fax: +49-30-2093 5934, home page http://www.wiwi.hu-berlin.de/wpol/

1 Introduction

How well do we understand business cycles and growth? One benchmark answer is provided by the stochastic neoclassical growth theory or real business cycle theory. Here, growth is understood to result from the steady progress of productivity, while business cycles result from fluctuations of productivity around this trend. In this paper, we examine how well a rather simple benchmark real business cycle model fits the key facts.

The paper first documents some key empirical facts in section 2. It then examines what aspects of the data are "explained" by a fairly standard simple real business cycle model, and which ones are not in section 3. Finally, concluding remarks are offered in section 4 on where this literature might be going and might have to go. In summary, the fit of the model at business cycle frequencies is quite compelling, given the simplicity of the model. However, the model does not seem to come much closer to the data than a "primitive" model, which just rescales TFP fluctuations. Moreover, the fit becomes less convincing, when no filtering is done. We conclude from this, that comparing theories to the data in this way is a demanding test. The challenge is to construct better quantitative theories that address the deficiencies documented here.

One of the major success stories in macroeconomic research over the last two decades has been the development of quantitative theories. The goal of these theories is to match empirical facts, while self-imposing two sources of discipline. The first discipline is that the theories ought to be based on sound microeconomic foundations. The theories need to spell out the environment as demanded in general equilibrium theory, e.g. Debreu (1959), i.e. they should specify preferences, technologies and endowments. This is in contrast to structural or nonstructural econometric models, which typically utilize only a subset of theoretical implications. The second discipline is that the comparison to the data ought to be quantitative. Statistics and numbers implied by the model are compared to the corresponding statistics and numbers in the data. This is in contrast to qualitative theorizing, which aims at providing insights into the general direction of economic effects without touching the issue of magnitudes.

There are now a plethora of quantitative theories available, and this paper is not the place to review that literature. In contrast to some perceptions, these theories allow for a wide range of possibilities such as many (and thus nonrepresentative) agents, sticky prices, increasing returns, monopolistic competition, indeterminacies, finite lives etc., see e.g. Farmer (1999). Some of these theories are constructed with the explicit aim to provide a fairly exhaustive understanding of the interplay of macroeconomic forces and for providing policy advice: the models by Christiano, Eichenbaum and Evans (2001b) for the US and subsequently Smets and Wouters (2002) for Europe are excellent examples.

Which of these theories will survive as leading explanations of observed phenomena will remain to be seen, and is subject to ongoing debate. Part of that debate is technical. First, these theories are typically not meant to explain "everything", leaving out certain aspects for the purpose of simplicity. Put differently, the theories are not meant to represent the true data generating processes. This makes usual econometric techniques such as maximum likelihood questionable for the purpose of estimation or data comparison.

The search for suitable techniques beyond calibration methods is on and has recently successfully focussed on a comparison of impulse responses both in the model as well as in VARs, see e.g. Christiano, Eichenbaum and Evans (2001a,2001b). Second, Occam's razor should be routinely applied, i.e. we should generally favor simpler rather than more complicated theories. Thus e.g., a model which only needs a "representative agent" and thus a single set of preferences to explain the same set of facts as a model using many agents should be preferred. To put it in pointed terms: Occam's razor is the opposite of realism. Reality is complicated. Our theories should not be, if they can avoid it.

A remarkably simple model is the benchmark real business cycle model by Hansen (1985). It is a full general equilibrium model in the sense of Debreu (1959), yet it can be stated in three lines and solved with a pocket calculator, performing nothing more spectacular than the solution of a quadratic equation, see Uhlig (1999). Finally, it is able to match observed statistics - fluctuations of GNP, hours worked, consumption and investment at business cycle frequencies - with remarkable quantitative accuracy. The model misses on a number of other dimensions, e.g. the volatility of real interest rates, see e.g. Hornstein and Uhlig (2000), but then again, it is not meant to explain all features of the data. Considerably richer models such as e.g. the models in Cooley (1995), Christiano, Eichenbaum and Evans (2001a,2001b) or Altig, David, Lawrence J. Christiano, Martin Eichenbaum and Jesper Linde (2002) go much further in matching a number of additional aspects. In particular, the literature has pointed to variable capacity utilization as a key component to improve further on the exercise performed here: see e.g. King and Rebelo (1999) for an exercise similar to ours, incorporating ariable capacity utilization.

In this paper, we stick to a version of Hansen (1985) model and reexamine its fit to the data in section 3. In order to do so and in order to make the comparison compatible with the empirical investigation to follow, we need to amend it with a few features such as accounting for population growth as well as for the difference between privately produced output and the sum of private consumption and investment. We confirm and demonstrate that it replicates certain business cycle features very well, but also confront the model to the raw, unfiltered data, where it does less well. We believe that this model or a model of this type is the benchmark to beat.

Real business cycle models have come under much attack recently: in fact, it has become fashionable to proclaim them fundamentally flawed as an explanation of business cycles. This may be so, but the challenge on the table is to provide alternative quantitative theories, which perform as least as well as the class of real business cycle models in explaining the data.

Key to the claim that productivity-shock driven theories are useful for understanding business cycles is the fact that productivity and labor are both procyclical. The most recent attack on the real business cycle paradigm thus comes from reexamining the Solow residual, i.e. total factor productivity, and demonstrating that it is not exogenous and therefore does not represent exogenous technological fluctuations. Galí (1999) and subsequently Francis and Ramey (2001) have therefore proposed to instead identify exogenous technology shocks by their long-run effects in a VAR. With their assumptions, they find that technology shocks actually lead to declines in labor rather than increases in labor

	output		hours		labor prod.		wages		TFP	
	HP	Δ	HP	Δ	HP	Δ	HP	Δ	HP	Δ
output	1	1								
hours	0.86	0.70	1	1						
labor prod.	0.54	0.68	0.04	-0.04	1	1				
wages	0.13	0.18	-0.11	-0.24	0.42	0.49	1	1		
TFP	0.81	0.85	0.40	0.23	0.92	0.96	0.32	0.40	1	1

Table 1: *Correlations, using the Hodrick-Prescott (HP)-filter and the first difference filter to detrend the data.*

input. If one accepts this conclusions, then technology shocks themselves cannot be a lead explanation of business cycles. The conclusion is subject to considerable debate, though, see e.g. Altig et al. (2002) and Uhlig (2002a,b). Whether other quantitative theories can rise to the challenge of replacing fluctuations in productivity as the engine of business cycles and provide a better quantitative fit remains to be seen.

2 Empirics: the facts to be explained

A remarkable feature of business cycle is that output and labor productivity are positively correlated. Any explanation of business cycles must be consistent with this fact. For example, a business cycle theory which is based on autonomous shifts in the labor supply due e.g. to leisure preferences, union demands or regulations or a theory which is based on shifts in demand for total output due e.g. to foreign trade or government policy will not replicate this observation without further refinements, since shifts along a given supply curve for output and assuming decreasing average productivity for labor along the aggregate production function will produce countercyclical, not procyclical labor productivity movements. This issue is well understood.

The facts are shown in table 1, using quarterly US data from 1947 to the fourth quarter of 2000[1]. The data are per person, age 16 or above. Let Y_H be the BLS output per hour in business, N be total hours in private business, W be nominal hourly compensation in business divided by the implicit price deflator for business, and P16 be total population age 16 or above. Let $LY_H = 100 * \log(Y_H)$, $LN = 100 * \log(N/\text{P16})$, $LW = \log(W)$ and construct the log of private business output per person LY as $LY = LY_H + LN$. The construction of Log TFP is explained below. Apply the HP-filter or first-difference this data to calculate the correlations above, using $\lambda = 1600$ for the HP filter.

A couple of results can be noticed. Output fluctuations are substantially positively correlated with fluctuations in both hours as well as labor productivity, but the correlation with wages is near zero. Labor productivity and hours are almost uncorrelated: this has been called the Dunlop-Tarshis observation, see Christiano-Eichenbaum (1992).

[1] I am grateful to Valerie Ramey for providing me with this data. The NIPA data are the same as the corresponding chain-weighted quarterly data available from the BEA. For the construction of TFP, see below

The table also contains correlations with total factor productivity, or TFP: more on its construction below. TFP is even more strongly procyclical than labor productivity. Moreover, and in contrast to labor productivity, TFP is positively correlated with hours worked. Real business cycle theories interpret the positive correlations between total factor productivity and output as indicating exogenous shifts in the production function over the cycle, triggering adjustment processes in all macroeconomic variables.

2.1 Capital and productivity: construction of the data

Total factor productivity is the statistic, which naturally emerges from growth accounting. Since capital shares and labor shares have remained fairly constant over time, one can try to represent total production in private businesses as a Cobb-Douglas production function, using e.g. the share ρ of capital income as exponent for capital,

$$Y_t = \Gamma_t K_{t-1}^{\rho} N_t^{1-\rho} \tag{1}$$

Here, y_t is output, k_{t-1} is capital, n_t is labor and Γ_t is TFP. To calculate Γ_t, one needs a number for ρ and one needs a time series for capital, which we shall now construct.

To that end, assume that capital evolves according to

$$K_t = (1-\delta)K_{t-1} + X_t \tag{2}$$

where X_t is gross investment.

With these two equations, values for δ and ρ and some auxilliary assumptions one can construct a time series for capital and for productivity. Observe that

$$\frac{K_t}{Y_t} = (1-\delta)\frac{Y_{t-1}}{Y_t}\frac{K_{t-1}}{Y_{t-1}} + \frac{X_t}{Y_t}$$

Assuming the ratios in this expression to be stationary and $\frac{Y_{t-1}}{Y_t}$ and $\frac{K_{t-1}}{Y_{t-1}}$ to be (approximately) uncorrelated, we find for the averages

$$\frac{\bar{K}}{Y} = \frac{1}{1-\frac{(1-\delta)}{g_Y}}\frac{\bar{X}}{Y}$$

where g_Y is the mean growth factor, i.e. the average of $\frac{Y_t}{Y_{t-1}}$. Given a level of depreciation δ and the average investment-output ratio, one can calculate the average capital-output ratio. Multiplying this ratio with initial output provides a guess for initial capital. With equation (2), one now obtains a time series for capital. We use the BEA NIPA real annual data starting in 1929 for gross investment and output and this calculation to obtain a time series for capital prior to 1947 and thereby a starting point for capital in 1947, used to construct the remaining data and comparison. Care needs to be taken in that the NIPA tables use annualized numbers also for the quarterly time series.

For the calculation prior to 1947, the investment series has been scaled up by the average ratio of investment plus durable consumption to investment, based on quarterly

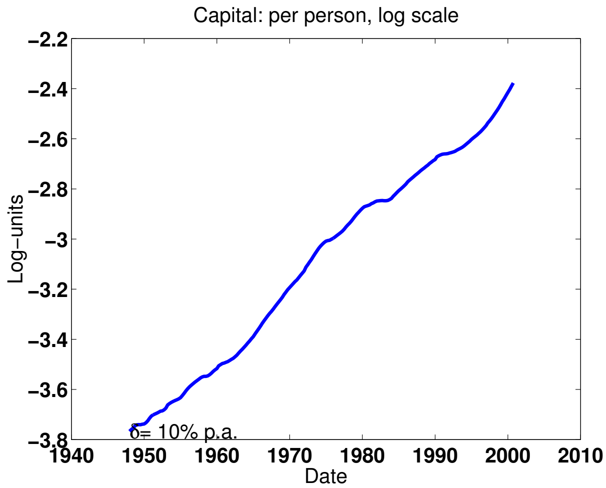

Figure 1: *Capital per person. The construction is explained in the text*

data after 1947. To calculate gross investment from NIPA data after 1947 and also for the comparisons to the theory below, we add durable consumption to gross private domestic investment. To obtain a time series for real durable consumption starting in 1947, we splice the real durable consumption series the same way the NIPA tables splice e.g. data on gross private domestic investment for obtaining real data starting in 1947. Similarly and for later comparison to the theory, we construct a quarterly time series for real consumption of nondurables and services, starting in 1947, which we simply interpret as consumption.

To obtain total factor productivity, calculate

$$\log \Gamma_t = \log Y_t - (\rho \log K_{t-1} + (1-\rho)\log N_t)$$

Two parameters, ρ and δ, are necessary for performing this calculation. We use $\rho = 0.30$, which is in the range of numbers typically used in the literature, and can be fixed from the share of capital income, and $\delta = 10\%$ on an annual basis (i.e. $\delta = 0.025$ for quarterly data), which again is a commonly used number. Furthermore, these parameters imply (together with the other parameters we choose further below), that the theory steady state investment-to-output ratio is 23%, a value very close to the corresponding average in the data, which is 24%. These numbers imply a steady state output-to-capital ratio of 12.7% for y_t/k_t.

2.2 Observations about capital, productivity and others.

Figure 1 shows the calculated capital time series per person over the age of 16. The time series and the results to follow are somewhat sensitive to the assumption about initial capital in 1947 - which we have constructed here based on annual data starting in 1929 - as well as the parameters.

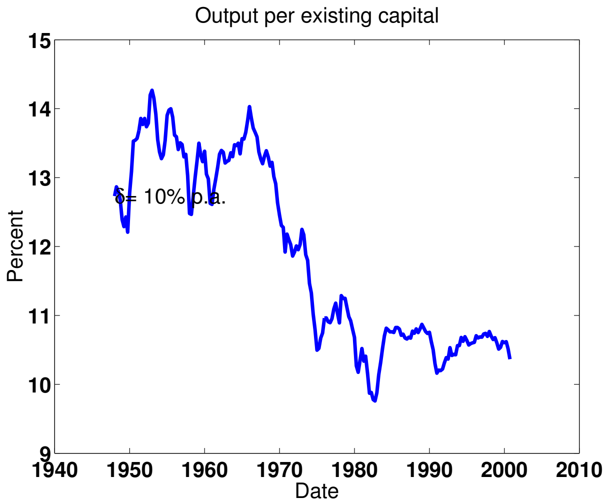

Figure 2: *Output per unit of capital. Note that the ratio is high before 1970 and low after 1970, indicating a possible break around 1970.*

Figure 2 shows the output-to-capital ratio, using per-capita data, i.e. it shows $(Y_t/\mathrm{P16}_t)/(K_{t-1}/\mathrm{P16}_{t-1})$. The ratio is higher than the steady state value of $g_{Y/\mathrm{P16}} *$ 12.7% prior to 1970, and falls below it afterwards, where $g_{Y/\mathrm{P16}}$ is the growth factor for $Y_t/\mathrm{P16}_t$. This probably indicates that one should really treat the sample has having a break somewhere around 1970 rather than assume a constant output-to-capital ratio throughout, around which the data fluctuates. We have not pursued a calculation, using a break: instead, we shall see that this break will occur in a number of pictures.

Figure 3 shows the implied time series for total factor productivity. TFP appears to be growing at a fairly constant pace over the entire sample, with some slowdown after 1970.

To examine this more closely, it is useful to study the growth rate of TFP and to compare it to the growth rate of labor productivity. To understand this relationship, note that (1) implies

$$(1-\rho)\log\frac{Y_t}{N_t} = \log\Gamma_t - \rho\log\frac{Y_t}{K_{t-1}}$$

relating labor productivity on the left hand side to TFP and the output-capital-ratio on the right hand side. This equation implies

$$g_{\Gamma_t} - (1-\rho)g_{\frac{Y_t}{N_t}} = \rho g_{\frac{Y_t}{K_{t-1}}} \tag{3}$$

i.e. the difference between the growth rate of TFP and $(1-\rho)$ times the growth rate of labor productivity is equal to ρ times the growth rate of the output-capital ratio shown in figure 2. Figure 4 compares the trends in labor productivity growth multiplied with $(1-\rho)$ and total factor productivity growth. The trends are constructed with the

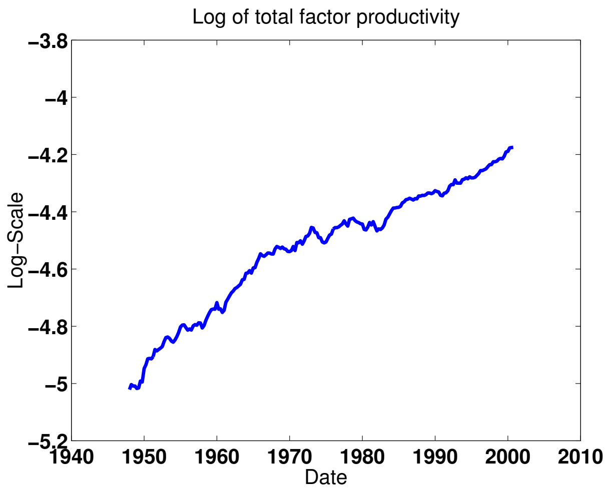

Figure 3: *Total factor productivity, using the calculated capital time series,* $\rho = 0.3$ *and* $\delta = 0.025$ *for quarterly data.*

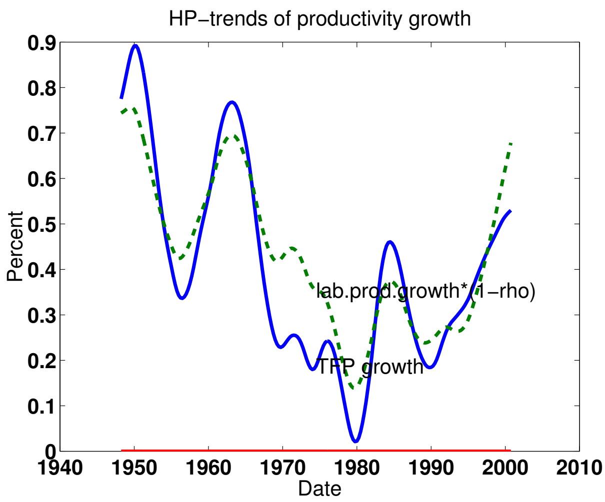

Figure 4: *Comparing trends in TFP and labor productivity. TFP is the solid line, whereas labor productivity is the dashed line. The trends were calculated with the HP filter. The trends are broadly consistent and the productivity slowdown from approximately 1970 to 1995 is clearly visible.*

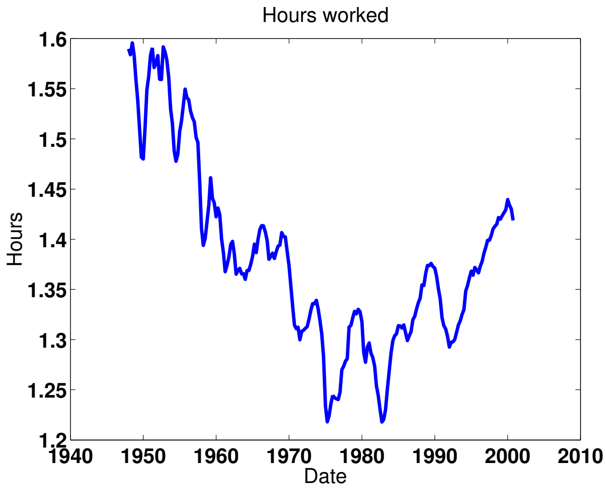

Figure 5: *Hours worked per person above the age of 16 (ignoring a scale factor).*

Hodrick-Prescott (HP) filter. The two trends are broadly consistent, with the differences "explained" by the changes in the output-capital ratio, see figure 2.

There are fluctuations in the broad trends in hours worked too, which can be glanced from figure 5. Hours per person are trending down until 1980 and then trending up again. Whether this time series is stationary or not underlies some of the recent debate about the role of productivity shocks, see Francis and Ramey (2001).

Examining the deviations from the HP-trend for output and juxtaposing it to the corresponding deviations in TFP is done in figure 6, visually reconfirming the results of table 1: TFP (and likewise labor productivity) comoves with output over the cycle.

There are two remaining bits to clean up before proceeding to the theory and to a quantitative comparison to the data: these issues are occasionally swept under the rug as of minor importance, but need to be treated carefully in order to allow for a clean comparison.

The first issue is population growth. We have concentrated on per-person data throughout. This makes sense: what matters for preferences is arguably not the total number of hours worked but the hours per person, and not total consumption but consumption per person. Figure 7 shows the annualized growth rates of the population above the age of 16 exhibit considerable fluctuations. These fluctuations matter e.g. for calculating the return to capital and a host of other issues: we thus need to incorporate them into the theory.

The second issue is that private consumption plus private investment does not equal private output, as simple business cycle models assume. The most appropriate method to deal with this issue is to also include a government sector as well as international trade into the model, at the price of additional complexity and the need to make choices. The alternative - slightly worse, but not much - is to treat the difference as due to exogenous

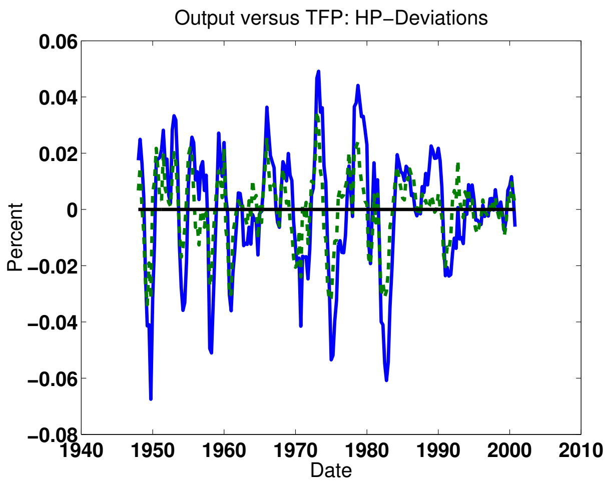

Figure 6: *Comparing business cycle frequency fluctuations in output (solid line) to those in TFP (dashed line). The fluctuations have been constructed, using the HP filter. Note the close correspondence of the two time series.*

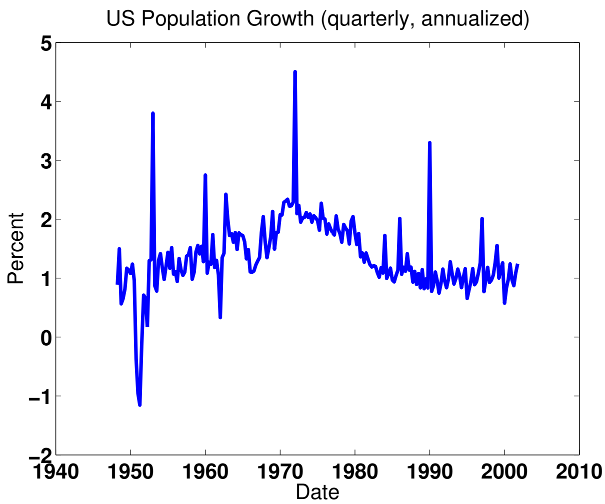

Figure 7: *Growth rates of the US population above the age 16.*

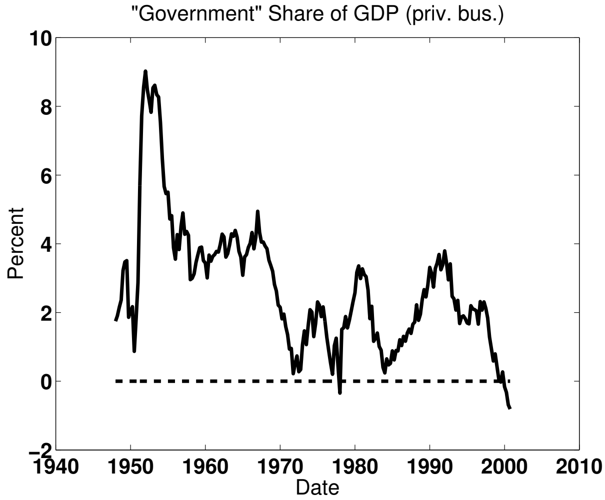

Figure 8: *"Government spending" as share of private business output per person. "Government spending" is constructed as the difference between private business output and private consumption plus private investment.*

government consumption. Figure 8 plots the difference between private business output and private consumption plus private investment. Examining a similar plot for actual government spending divided by private business output shows that the difference between these two is fairly constant, which in turn roughly corresponds to the difference between private business output and total GDP. We therefore feel that we do not make a major mistake with this interpretation, and have instead provided for a convenient way to interpret the difference.

3 A real business cycle model

What can explain these (and other) macroeconomic facts? Real business cycle models supply one reasonably successful set of theories. We shall examine one prototype example, building on Hansens (1985) benchmark real business cycle model. We amend his model to allow for stochastic population growth and for stochastic and exogenous government spending for the reasons described towards the end of the previous section.

We shall provide a comparison of the model time series to the data time series for the period 1960 to 2001 in order to minimize potential "beginning-of-sample" effects which may distort our calculations of the capital stock and TFP for 1947.

We use the following notational conventions. Capital letters, such as K_t denote economy-wide aggregates. Small letters with a caret such as $\hat{k}_t$ are per person above the age of 16, i.e., $\hat{k}_t = K_t/\mathrm{P16}_t$. Small letters with a tilde such as $\tilde{k}_t$ are detrended to achieve stationarity: exactly how is explained further below. Finally, k_t itself will denote the log-deviation of $\tilde{k}_t$ from its steady state $\bar{k}$, $k_t = \log \tilde{k}_t - \log \bar{k}$. K_t denotes capital available at the end of period t and used for production at the beginning of period $t+1$, C_t is consumption in period t, N_t is hours worked, Y_t is output, G_t is government spending, $\mathrm{P16}_t$ is the population aged 16 or above. Values without time subscripts or bars denote parameters. $E_t[\cdot]$ is the conditional expectation at date t.

We assume that there is a representative household, who maximizes its infinite-horizon expected discounted utility subject to the resource constraint of the economy,

$$\begin{aligned} &E\left[\sum_{t=0}^{\infty} \beta^t \left(\log \hat{c}_t - A\hat{n}_t\right)\right] \\ &\text{s.t.} \quad C_t + G_t + K_t = Y_t + (1-\delta)K_{t-1} \\ &\qquad\;\; Y_t = \Gamma_t K_{t-1}^{\rho} N_t^{1-\rho} \end{aligned}$$

and given K_{-1}, where e.g. $C_t = \hat{c}_t * \mathrm{P16}_t$ and where Γ_t, $\mathrm{P16}_t$ and G_t are exogenous stochastic processes. A few remarks are in order. First, faster population growth will ceteris paribus decrease utility: with a given stock of total capital and since there are decreasing returns to labor, a larger population will result in less output per person $\hat{y}_t = Y_t/\mathrm{P16}_t$. An alternative would be to multiply the per-period felicity function $\log \hat{c}_t - A\hat{n}_t$ with the size of the population $\mathrm{P16}_t$, but this would imply that the total labor supply, not the per-person labor supply would be stationary, an implication incompatible with the data and necessitating another utility function. Finally, the linearity in the disutility for hours worked implies a rather high elasticity of labor supply. This can be made consistent with the micro-evidence of individually low elasticity of labor supply with respect to wages by incorporating indivisibility of labor and using the lottery formulation to aggregate individual labor supplies, see Hansen (1985) or the original source, Rogerson (1988) for details.

The model does not have much to say about real wages, since nothing is said about the way labor markets operate. One can decentralize the model above as a competitive equilibrium and assume that wages are equal to market-clearing spot wages for labor, in which case there is no difference between real wages and labor productivity. But even in a competitive equilibrium, there is no reason for measured wages to correspond to these spot wages. For example, it is perfectly consistent with this model and with a competitive equilibrium, that agents specify labor contracts involving labor supply for several periods with a variety of termination options and insurance possibilities. Viewed from an asset pricing perspective, labor contracts can be understood as derivative securities based on the fundamental underlying assets such as the spot market for labor and the spot market for capital. What we observe as wage then corresponds to the price for these derivative securities. Without further specifying the exact nature of these contracts, it is impossible to make any meaningful statement about the behaviour of real wages: therefore, we do not.

3.1 Specifying the details and solving the model

The logarithm of population and the logarithm of TFP may be usefully regarded as random walks with drift: the AR(1) coefficients are 0.9997 and 0.9931 respectively, the AR(1) coefficients of their log-differences are 0.5803 and 0.0442. An alternative is to regard the logarithm of population and the logarithm of TFP as trend-stationary. In either case, one can reformulate the model in terms of stationary variables by detrending all variables except labor by $Z_t = \Gamma_t^{1/(1-\rho)}$, which can be interpreted as labor productivity (up to a constant factor), if the output-capital ratio was constant, see equation (3). I.e., let $\tilde{k}_t = \hat{k}_t / Z_t = K_t/(Z_t * \mathrm{P16}_t)$, $\tilde{c}_t = \hat{c}_t/Z_t = C_t/(Z_t * \mathrm{P16}_t)$, $\tilde{y}_t = \hat{y}_t/Z_t$, $\tilde{g}_t = \hat{g}_t/Z_t$, $\tilde{n}_t = \hat{n}_t$. The model can then be equivalently rewritten as

$$E\left[\sum_{t=0}^{\infty} \beta^t \left(\log Z_t + \log \tilde{c}_t - A\tilde{n}_t\right)\right]$$

$$\text{s.t.} \quad \tilde{c}_t + \tilde{g}_t + \tilde{k}_t = \tilde{y}_t + (1-\delta)\left(\frac{\mathrm{P16}_t Z_t}{\mathrm{P16}_{t-1} Z_{t-1}}\right)^{-1} \tilde{k}_{t-1}$$

$$\tilde{y}_t = \left(\frac{\mathrm{P16}_t Z_t}{\mathrm{P16}_{t-1} Z_{t-1}}\right)^{-\rho} \tilde{k}_{t-1}^{\rho} \tilde{n}_t^{1-\rho}$$

Note that population growth and shocks to population growth play the same role as TFP growth and shocks to TFP growth: a difference comes about in the dynamics only because population growth and TFP growth have different degrees of persistence. Also, an increase in TFP seems to lead to a decrease in output $\tilde{y}_t$, holding the inputs $\tilde{k}$, $\tilde{n}$ constant. But remember, that $\tilde{y}_t$ is itself detrended: the level of output will increase with a TFP shock, just as before.

Substituting $\tilde{y}_t$ in the feasibility constraint with the help of the production function and writing $\tilde{\lambda}_t$ for the Lagrange multiplier for the resulting equation, one now easily derives the first order conditions

$$\begin{aligned} \tilde{\lambda}_t &= \frac{1}{\tilde{c}_t} \\ A &= \tilde{\lambda}_t (1-\rho) \frac{\tilde{y}_t}{\tilde{n}_t} \\ \tilde{\lambda}_t &= \beta E_t[\tilde{\lambda}_{t+1} \tilde{R}_{t+1}] \end{aligned}$$

where

$$\tilde{R}_t = \rho \frac{\tilde{y}_t}{\tilde{k}_{t-1}} + (1-\delta)\left(\frac{\mathrm{P16}_t Z_t}{\mathrm{P16}_{t-1} Z_{t-1}}\right)^{-1}$$

can be read as the return to investing in capital. It is related to the return

$$R_t = \rho \frac{Y_t}{K_{t-1}} + 1 - \delta = \left(\frac{\mathrm{P16}_t Z_t}{\mathrm{P16}_{t-1} Z_{t-1}}\right) \tilde{R}_t$$

in the first formulation of the model.

In order to solve the model and analyze its quantitative properties, one needs values for the parameters. We have taken fairly standard values: $\rho = 0.30$, $\delta = 0.025$, $\beta = 1/1.01$, $\bar{n} = 1.34$, and thus $A = 0.71$. The parameter A has been chosen so that the steady state supply of hours per person equals the average labor supply in the data for the period 1960 to 2001. The exogenous processes $\tilde{\gamma}_t = \log \Gamma_t - \log \Gamma_{t-1}$, $\tilde{\nu}_t = \log P_t - \log P_{t-1}$ and $\tilde{g}_t$, have been assumed to be AR(1) processes with the means and AR(1) coefficients ψ_γ, ψ_ν and ψ_g respectively given by $\bar{\gamma} = \bar{z}(1-\rho) = 0.40\%$, $\psi_\gamma = \psi_z = 0.044$, $\bar{\nu} = 0.34\%$, $\psi_\nu = 0.58$, $\bar{g} = 0.095$, (thus, $\bar{g}/\bar{y} = 0.03$) and $\psi_{\tilde{g}} = 0.97$. We have not taken the log of $\tilde{g}_t$, since $\tilde{g}_t$ occasionally takes negative values. Government spending $\tilde{g}_t$ is a very persistent process: its AR(1) coefficient is near unity, and if it were to play a larger role, it may be sensible to use a better fitting regression to describe its stochastic properties. In any case, one probably does not make much of a mistake in assuming simple AR(1) processes: the main difference to a more elaborate specification is that agents assume the processes to be somewhat less predictable than they actually are.

Given these model parameters, one can easily calculate the steady state $\bar{y}, \bar{k}, \bar{n}, \bar{c}$ and $\bar{R}$ for $\tilde{y}_t, \tilde{k}_t, \tilde{n}_t, \tilde{c}_t$ and $\tilde{R}_t$. To calculate the dynamics around the steady state, consider the log-deviations

$$\begin{aligned} c_t &= \log \tilde{c}_t - \log \bar{c} \\ r_t &= \log \tilde{R}_t - \log \bar{R} \end{aligned}$$

etc., with the exception that $g_t = \tilde{g}_t - \bar{g}$ to allow for the possibility that it can take negative values. Likewise, let

$$\begin{aligned} \gamma_t &= \tilde{\gamma}_t - \bar{\gamma} \\ &= (\log \Gamma_t - \log \Gamma_{t-1}) - \text{mean}(\log \Gamma_t - \log \Gamma_{t-1}) \\ &= (1-\rho)\left(\log\left(\frac{Z_t}{Z_{t-1}}\right) - \text{mean}\left(\log\left(\frac{Z_t}{Z_{t-1}}\right)\right)\right) \\ \nu_t &= \tilde{\nu}_t - \bar{\nu} \\ &= (\log \mathrm{P16}_t - \log \mathrm{P16}_{t-1}) - \text{mean}(\log \mathrm{P16}_t - \log \mathrm{P16}_{t-1}) \end{aligned}$$

With that, the exogenous AR(1) processes can be written as

$$\begin{aligned} \gamma_t &= \psi_\gamma \gamma_{t-1} + \epsilon_{\gamma,t} \\ \nu_t &= \psi_\nu \nu_{t-1} + \epsilon_{\nu,t} \\ g_t &= \psi_g g_{t-1} + \epsilon_{g,t} \end{aligned}$$

and the residuals $\epsilon_{\gamma,t}, \epsilon_{\nu,t}$ and $\epsilon_{g,t}$ can be calculated from the data.

The equations characterizing the solution of the model, i.e. the feasibility constraint, the production function and the first order conditions, can be log-linearized to deliver a linear system in the log-deviations. The linear system can then be solved for the stable recursive equilibrium law of motion in these log-deviations. The calculations do not involve anything more complicated than the solution of a quadratic equation, and can reasonably be performed with pencil, paper and pocket calculator: details are contained in Uhlig (1999).

The result is a recursive equilibrium law of motion, relating the current date-t values for the log-deviations of all variables to the state variables, i.e. to the endogenous state variable k_{t-1} and the three exogenous state variables γ_t, ν_t and g_t. One obtains

$$\begin{aligned}
k_t &= 0.91k_{t-1} - 1.31\gamma_t - 1.13\nu_t + 0.02g_t \\
c_t &= 0.46k_{t-1} - 0.62\gamma_t + 0.04\nu_t - 0.15g_t \\
y_t &= -0.06k_{t-1} + 0.03\gamma_t - 1.09\nu_t + 0.36g_t \\
n_t &= -0.52k_{t-1} + 0.65\gamma_t - 1.13\nu_t + 0.51g_t \\
r_t &= -0.05k_{t-1} - 1.37\gamma_t - 1.00\nu_t + 0.02g_t
\end{aligned}$$

and $\lambda_t = -c_t$. The initial values can be computed from the data as

$$k_{-1} = -8.3\%, \gamma_0 = 1.91\%, \nu_0 = 0.34\%, g_0 = 0.021$$

with $t = 0$ corresponding to the first quarter of 1960. With the equations above, these initial values and the data for γ_t, ν_t and g_t, one can easily construct impulse responses or simulate the model by recursively calculating the values for k_t, c_t, y_t, n_t, r_t for $t = 0, \ldots, T$. By adding the logarithm of the steady state and adding $\log Z_t$ (except for n_t), one can obtain simulations comparable to the log of the actual data, and compare.

It is worth emphasizing, that this model really is remarkably simple. It is stated in three lines. The solutions involves nothing more than solving a quadratic equation. And the model can be simulated, using a few linear equations stated above.

3.2 Results

To get a sense for the dynamics and for the comparison with the empirical results in the next section, it is useful to examine impulse response functions to a productivity growth shock $\epsilon_{\gamma,0} = 1\%$. Rather than plot the response of the TFP-detrended values y_t, c_t, etc., we show the response of the non-TFP-detrended variables, i.e. we show $y_t + \zeta_t$, $c_t + \zeta_t$, etc., where

$$\zeta_t = \frac{1}{1-\rho}\left(\sum_{s=0}^{t} {\psi_\gamma}^s\right)\epsilon_{\gamma,0} = \frac{1-{\psi_\gamma}^{t+1}}{(1-\rho)(1-\psi_\gamma)}\epsilon_{\gamma,0}$$

is the change in the detrending factor $\log(\Gamma_t^{1/(1-\rho)})$ due to the shock to productivity growth, $\epsilon_{\gamma,0} = 1\%$. These non-TFP-detrended variables show, how the per capita level variables change in percent due to the productivity growth shock.

The result can be seen in figure 9. As is typical in this model, output, TFP and hours worked all rise simultaneously. Labor then slowly declines after the initial jump upwards, while consumption and capital show a hump-shaped response, following an initial jump upwards. These are intuitively appealing results, and they are good news. We know from table 1, that these variables are positively correlated. The model provides dynamic

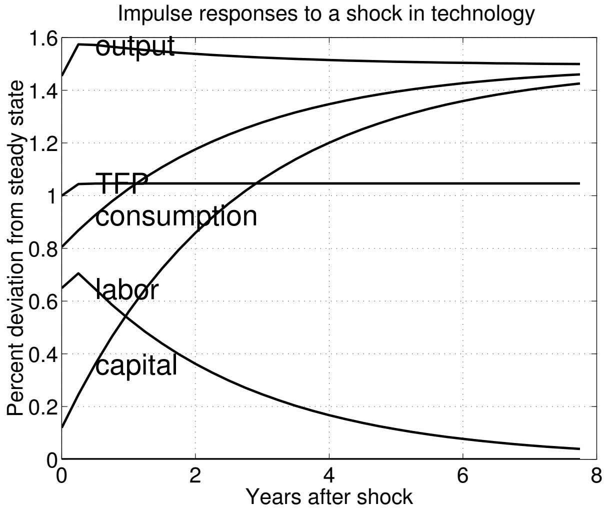

Figure 9: *Impulse responses to a technology shock, i.e. a shock to total factor productivity growth.*

responses consistent with these facts. Investigating this dynamic response is at the heart of the next section.

Figure 10 shows the response to a surprise one percent increase in the population above the age of 16: this could e.g. be immigration or (by reversing all signs) the response to a violent disease which wipes out 1% of the population. Recall that all the variables are in per capita terms: thus, initially, output and capital per capita decrease by approximately as much in percentage terms as the population increases. Consumption, interestingly, does not do so: apparently, the shortfall is made up by "consuming" the capital stock, i.e., by initially reducing investment.

The model is meant to explain the business cycle movements in the data. Thus, it is sensible to compare the data to a simulation of the model, where the actual data for TFP, population growth and "government spending" has been used as input. We shall consider the logarithm of all variables, and first compare the deviations from an HP trend.

The results are in table 2. The dotted lines indicate the data: they are very close to the solid lines, coming out of the simulations. This is a success. They show that this simple model can amazingly convincingly replicate the business cycle features of the data. Any other alternative model needs to be able to provide pictures which provide a similarly good fit.

This can also be seen from the raw statistics in table 3. The theory is able to replicate the size of the fluctuations, as indicated by the standard deviations, and comes close to the correlations observed in the data. While there are population shocks and government spending shocks in this model, this is essentially a one-shock model, driven mainly by productivity shocks. Enriching the model with additional sources of shocks might enable it to come closer to matching the observed correlations. The purpose here is to investigate,

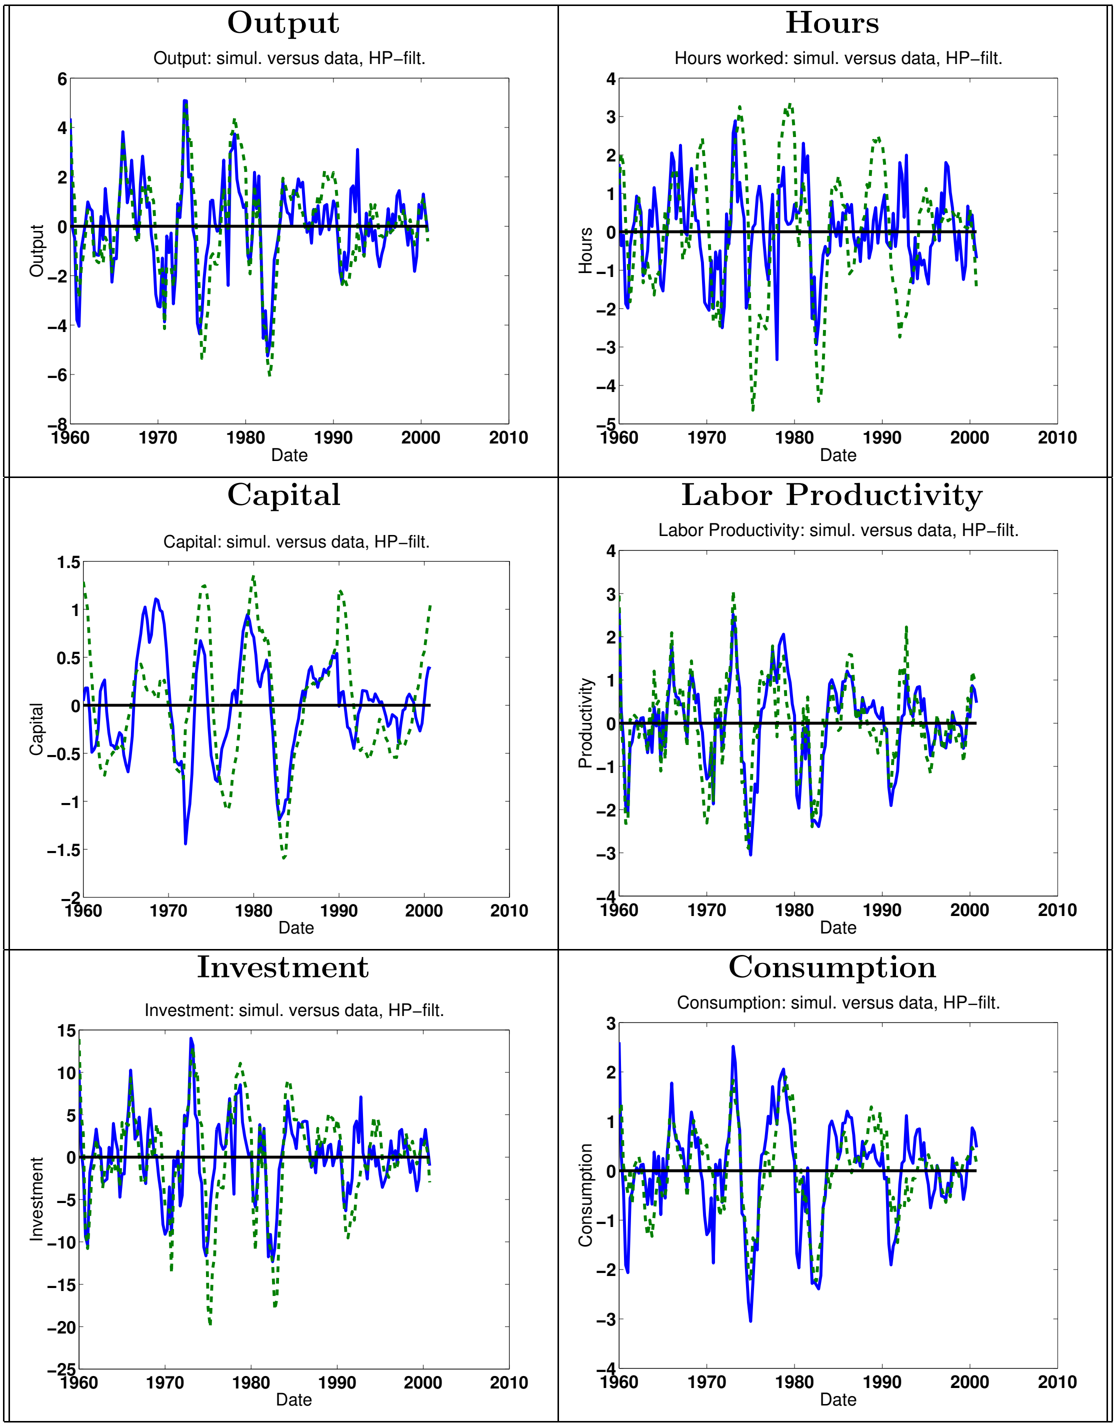

Table 2: *Comparison of the business cycle component in the simulations (solid line) and in the data (dashed line). The lines correspond very closely.*

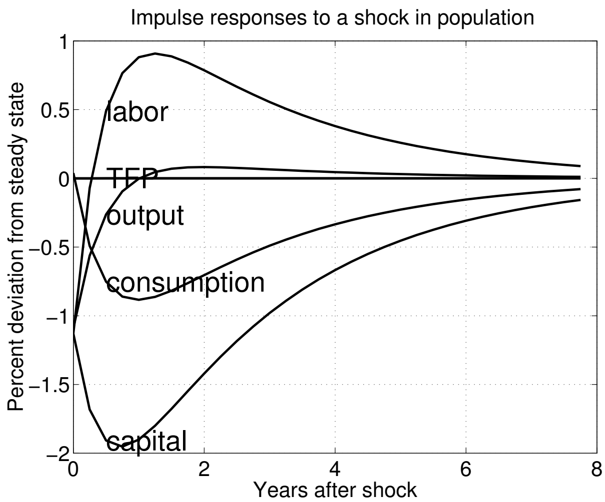

Figure 10: *Impulse responses to a population shock, i.e. a shock to population growth.*

	output	TFP	labor prod.	hours	cons.	inv.	cap.
std.dev: data	2.01	1.22	1.05	1.67	0.83	5.95	0.64
std.dev: theory	**1.83**	**1.22**	**1.03**	**1.09**	**1.03**	**4.58**	**0.52**
Correlations							
TFP: data	0.81	1.00					
theory	**0.94**	**1.00**					
lab.prod.	0.56	0.93	1.00				
	0.85	**0.97**	**1.00**				
hours	0.85	0.39	0.05	1.00			
	0.87	**0.66**	**0.47**	**1.00**			
cons.	0.83	0.66	0.45	0.72	1.00		
	0.85	**0.97**	**1.00**	**0.47**	**1.00**		
inv.	0.94	0.81	0.56	0.78	0.72	1.00	
	0.97	**0.97**	**0.87**	**0.80**	**0.87**	**1.00**	
capital	0.40	-0.03	-0.13	0.57	0.30	0.24	1.00
	0.40	**0.26**	**0.34**	**0.34**	**0.34**	**0.22**	**1.00**

Table 3: *Data-theory comparison of standard deviations and correlations with output for the business cycle frequencies for the time series listed. The business cycle component was calculated, using the HP-filter.*

data series	with theory series	with data TFP
TFP	1.00	1.00
output	0.81	0.81
labor productivity	0.86	0.93
hours	0.30	0.39
consumption	0.74	0.66
investment	0.73	0.81
capital	0.58	-0.03

Table 4: *Comparing the correlations of the results at business cycle frequencies. The second column shows the correlation of the data time series with the corresponding simulated time series, e.g. output in the data with output in the model. The third column shows the correlation of the data time series with the data TFP series, e.g. output in the data with TFP in the data. Note that there is generally not much improvement, when moving from the correlation with TFP to the correlation with the corresponding theory time series. All time series have been HP-filtered to compare the results at business cycle frequencies.*

in how far this simple model can already do the trick. The answer is: it does pretty well.

But can one really proclaim success, based on these results? Further investigations cast doubt on this conclusion. Take the results in table 4. There, the correlation between the data series and the simulated series have been juxtaposed to the correlations between the data series and the data on TFP: all series again are in logs and have been HP-filtered. Except for the time series on capital, the theory does not deliver a closer fit to the data than the raw TFP series alone, possibly appropriately rescaled. This is bad news, considering that the TFP series has been used as input into the simulations.

To put it differently, let e.g. $y_{HP,t}$ indicate the business cycle component of $\log(\hat{y}_t)$, obtained with the HP-filter. A model which simply appropriately rescales the business cycle component of TFP by postulating

$$\begin{aligned} y_{HP,t} &= \frac{\text{std.dev}(y_{HP,t})}{\text{std.dev}(TFP_{HP,t})} TFP_{HP,t} \\ c_{HP,t} &= \frac{\text{std.dev}(c_{HP,t})}{\text{std.dev}(TFP_{HP,t})} TFP_{HP,t} \end{aligned}$$

etc. does just as well as the real business cycle model above in "explaining" the data, except for capital. The propagation mechanisms in the model do not add anything towards our understanding of the model. This is essentially the point made by Cogley and Nason (1993, 1995): the internal propagation mechanism of the model is weak, the seemingly excellent fit shown in e.g. the figures of table 2 mainly come about due to the dynamics of TFP, which the model takes as given.

So, while there is an excellent fit, this is not in itself good news. It simply shows that the TFP fluctuations are essential for understanding business cycle fluctuations, but it does not show that the responses postulated by real business cycle models are. What the business cycle model does achieve is the appropriate scaling, i.e., it does explain why e.g.

investment fluctuates strongly and consumption weakly over the cycle.

One can take the comparison a step further and investigate the quality of the fit between simulations and data without applying the HP-filter. This is perhaps an unfair comparison, as the real business cycle models were not meant to explain the low frequency movements in the data as well. But since the model builds on the neoclassical growth model, and since eventually one wishes to understand these low-frequency dynamics as well, this is a sensible issue to investigate.

The results are shown in table 5. While there is some agreement between these lines, there are uncomfortable gaps. The model is often able to match the general trend as well as fluctuations at business cycle frequencies, but misses swings at somewhat longer frequencies. Given the insights above, this may not be regarded as too surprising. We have already seen that there is a close match at business cycle frequencies. Furthermore, per construction, $Z_t = \Gamma_t^{1/(1-\rho)}$ is used for detrending or re-trending all time series except hours worked. I.e., the general trend may again be nothing more than TFP plus parameter choices in the theory to match the investment-to-output ratio.

One might therefore suspect that detrending these series with Z_t and comparing the unfiltered version will reveal uncomfortable discrepancies. This is indeed the case, as table 6 reveals. While there are fits at business cycle frequencies, and while the fit e.g. for consumption looks fine, there is considerable disagreement at lower frequencies e.g. for capital, output or investment. It is interesting, though, that the lower frequency movements in hours worked are matched reasonably well.

These gaps at lower frequencies and likewise the break visible in the output-to-capital ratio in figure 2 can perhaps be fixed with slow-moving changes in depreciations rates or slow swings in TFP growth, as indicated by figure 4. The real business cycle model was not meant to address these facts. But given that it apparently does not do much more than rescale TFP growth fluctuations at business cycle frequencies and TFP trends for the long-run behavior, one would have hoped for better news here.

In sum, the real business cycle model apparently does contain some ingredients, which are there to stay - the emphasis on TFP, the desire for consumption smoothing - but lots of the details can be subject to a complete overhaul without doing much harm to the quality of the fit. The research agenda for finding appropriate models of the business cycle is, again, wide open. That does not mean that anything goes. The standards of comparison applied above and in the literature are rather stringent, and need to be applied to any alternative. These standards have been successful in ruling out lots of explanations already. It remains to be seen, which theories will eventually stand the test of time as convincing and possibly alternative explanations of the cycle.

4 Where do we go from here? Concluding Remarks

Where do we go from here? The investigation above delivers several positive as well as negative conclusions. On the positive side, it has been shown that productivity movements - be they TFP movements, labor productivity movements or technology shocks identified

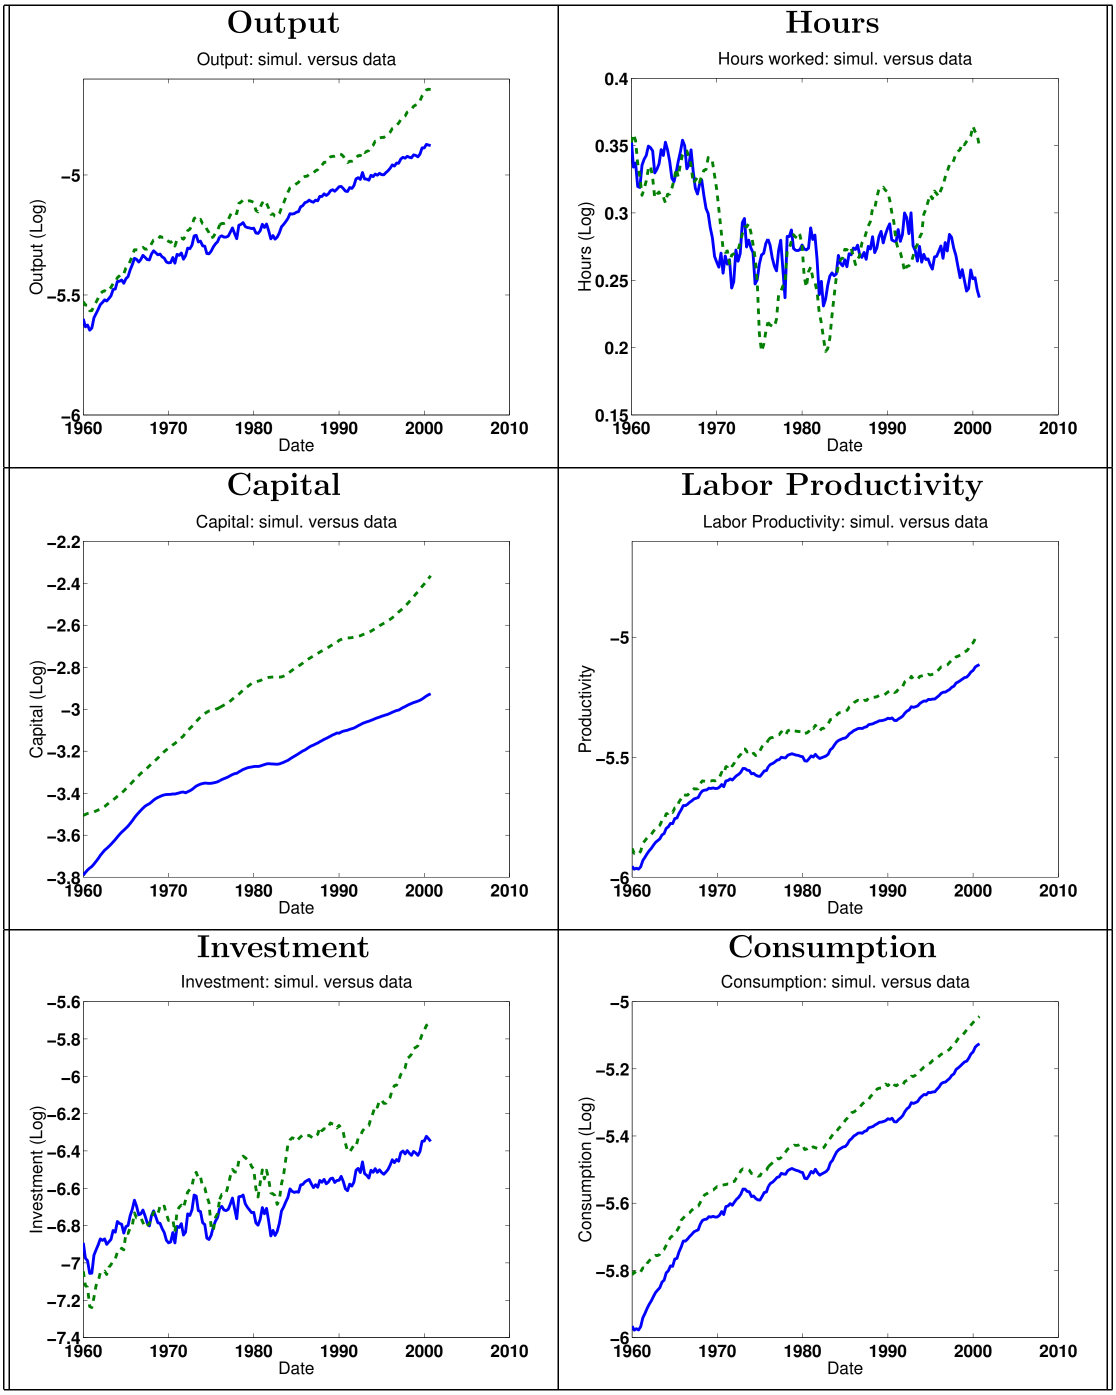

Table 5: *Comparison of the unfiltered time series in the simulations (solid line) and in the data (dashed line), using log scales. The agreement is close, but gaps remain.*

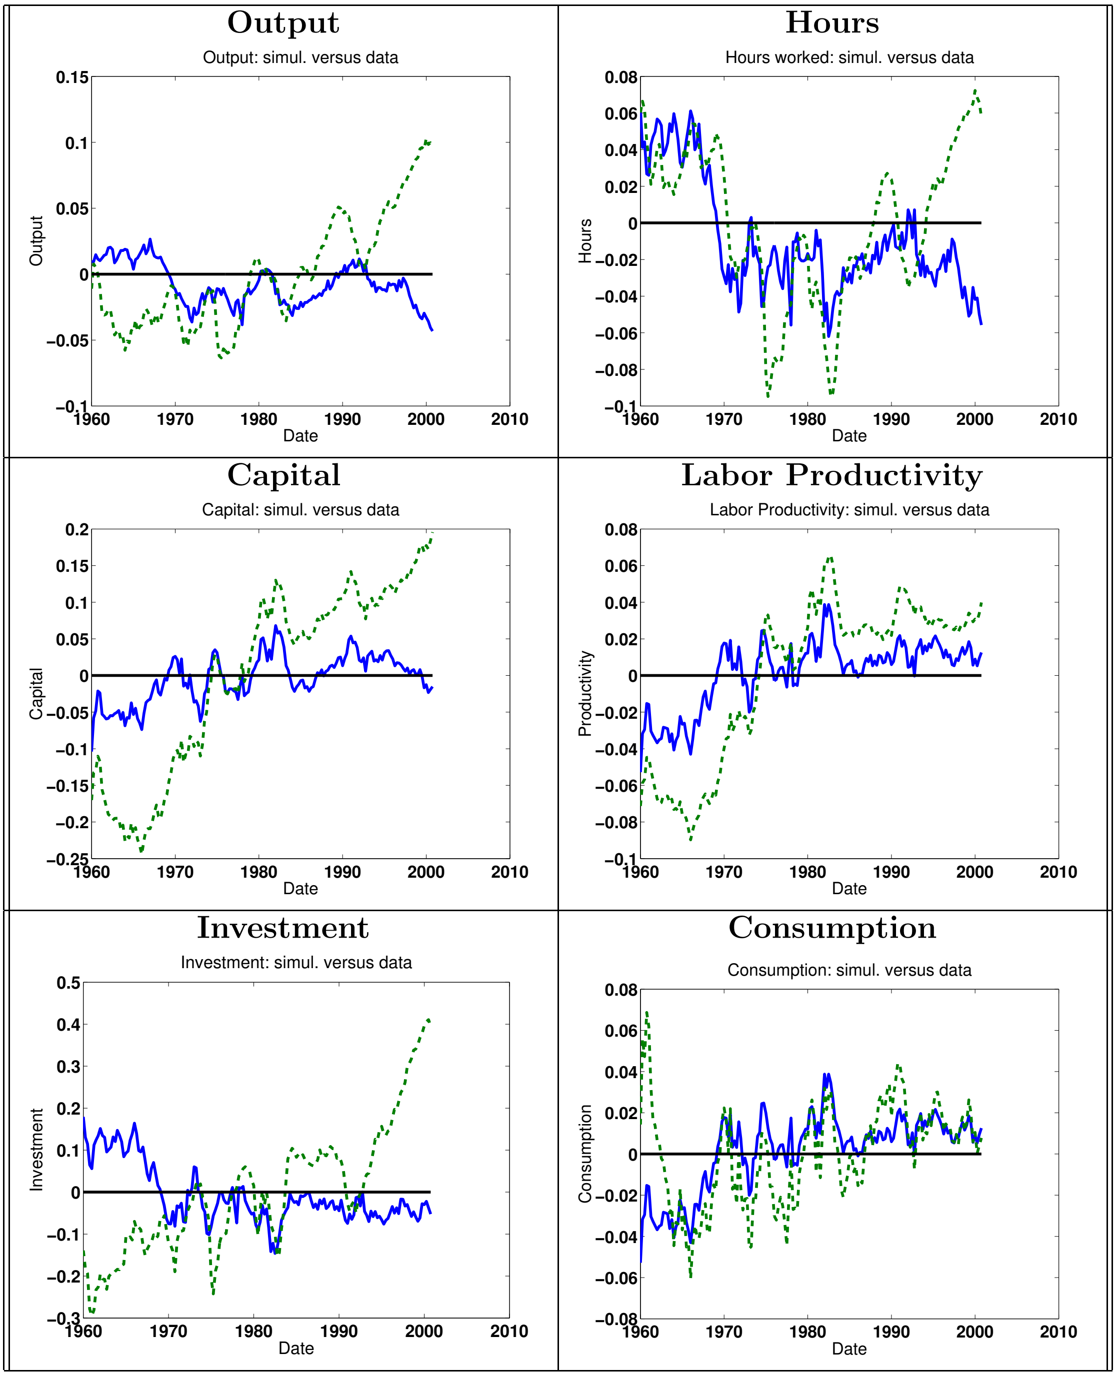

Table 6: *Comparison of the unfiltered time series in the simulations (solid line) and in the data (dashed line), using log scales. The productivity trend* $Z_t = \Gamma_t^{1/(1-\rho)}$ *has been removed. Uncomfortably big gaps are visible.*

in a reasonable manner - are key to understanding macroeconomic dynamics: they are the main source of movements in the medium-to-long run.

Whether these movements are helpful for understanding business cycle dynamics and whether technology shocks lead to rises or declines in hours worked is more of an open issue. The propagation mechanics of real business cycle theory is weak: the success of these theories in matching the facts is mainly due to the high correlation of TFP with a number of macroeconomic variables rather than a deepened understanding of the response to TFP movements. On the other hand, it may not be wise to already consider recent claims in the literature, that technology shocks lead to declines rather than rises in hours worked, as robust conclusions.

The literature has provided us with tools both in theory as well as in time series econometrics to probe further into these issues, and we must. Plant-level data should enable researchers to reach more conclusive results regarding the effects of technology shocks. Quantitative theories need to be build which go beyond replicating correlations with TFP and getting steady states and the size of the fluctuations right. The level of the debate has risen considerably. There is a large number of litmus tests we are able to apply now to separate good theories from bad or good and robust econometric conclusions from bad or nonrobust ones. These tests have been instrumental in discarding many ideas once constituting core macroeconomic teachings, but the search for convincing theories and econometric models is far from over. It remains to be seen, which theoretical and empirical models will survive these stringent tests: I am confident, that some will eventually.

References

[1] Altig, David, Lawrence J. Christiano, Martin Eichenbaum and Jesper Linde (2002), "Technology Shocks and Aggregate Fluctuations," draft, Northwestern University.

[2] Christiano, Lawrence J. and Martin Eichenbaum (1992), "Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations," American Economic Review; 82(3), June, pages 430-50.

[3] Christiano, Lawrence J., Martin Eichenbaum and Charles Evans (2001a), "Monetara Policy Shocks: What Have We Learned and to What End?," in J. Taylor and M. Woodford, *Handbook of Macroeconomics*, North-Holland, Amsterdam.

[4] Christiano, Lawrence J., Martin Eichenbaum and Charles Evans (2001b), "Nominal rigidities and the dynamic effects of a shock to monetary policy," draft, Northwestern University.

[5] Cogley, Timothy and James M. Nason (1993), "Impulse Dynamics and Propagation Mechanisms in a Real Business Cycle Model," Economics Letters; 43(1), 77-81.

[6] Cogley, Timothy and James M. Nason (1995), "Output Dynamics in Real-Business-Cycle Models," American Economic Review; 85(3), June, 492-511.

[7] Cooley, Thomas F. (1995), ed., *Frontiers of Business Cycle Research*, Princeton University Press, New Jersey.

[8] Debreu, Gerard (1959), "Theory of Value," New York: Wiley.

[9] Evans, Charles L. (1992), "Productivity Shocks and Real Business Cycles," Journal of Monetary Economics 29, April, 191-208.

[10] Farmer, Roger E.A. (1999), *The Macroeconomics of Self-fulfilling prophecies*, 2nd edition, MIT Press, Cambridge, Massachusetts.

[11] Francis, Neville and Valerie A. Ramey (2001), "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited," draft, UCSD.

[12] Galí, Jordi (1999), "Technology, Employment and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?", American Economic Review, vol. 89, no. 1, 249-271.

[13] Hansen, G.D. (1985), "Indivisible Labor and the Business Cycle," Journal of Monetary Economics 16, 309-27.

[14] Hornstein, Andreas and Harald Uhlig (2000), "What is the real story for interest rate volatility," German Economic Review, inaugural issue vol. 1, issue 1, February, 43-67.

[15] King, Robert G. and Sergio T. Rebelo (1999), "Resuscitating Real Business Cycles,", chapter 14 in John B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, Elsevier: Amsterdam, 927-1007.

[16] Rogerson, Richard (1988), "Indivisible Labor, Lotteries and Equilibrium," Journal of Monetary Economics; 21(1), January, pages 3-16.

[17] Smets, Frank and Raf Wouters (2002), "An estimated stochastic dynamic general equilibrium model of the Euro area," draft, European Central Bank.

[18] Uhlig, Harald (1999), "A toolkit for analysing nonlinear dynamic stochastic models easily," in Ramon Marimon and Andrew Scott, eds, *Computational Methods for the Study of Dynamic Economies*, Oxford University Press, Oxford, 30-61.

[19] Uhlig, Harald (2002a), "What moves real GNP?" draft, Humboldt University Berlin.

[20] Uhlig, Harald (2002b), "Do productivity shocks lead to a decline in labor?", draft, Humboldt University Berlin.