Can Modern Theories of Structural Change Fit Business Cycles Data?*

Loris Rubini[†] Alessio Moro[‡]

November 2024

Abstract

We investigate the ability of workhorse structural change models in accounting for the business cycle properties of an economy. We consider three different preferences specifications: Herrendorf, Rogerson and Valentinyi (2014, HRV), Boppart (2014), and Comin, Lashkari and Mestieri (2021, CLM), paired with standard sectoral production functions with random total factor productivity (TFP) shocks. In each case, we estimate preference parameters using long-run structural change data, and common TFP processes calibrated on observed relative prices. Our main results can be summarized by: i) all models display a volatility of aggregate variables substantially lower than the data, but they account for a large fraction of the volatility of consumption relative to GDP; ii) at the sectoral level, only CLM accounts for a substantial fraction of absolute and relative volatility; iii) all models do reasonably well in accounting for the cyclicality of aggregate GDP components; and iv) only HRV can account for the cyclicality of sectoral variables.

JEL Classification: E32, L16, O41.

Keywords: Structural Change, Stochastic Growth, Real Business Cycles.

*We thank Timo Boppart, Diego Comin, Berthold Herrendorf, Martí Mestieri and Galo Nuño for interesting discussions, and participants at the 2022 Midwest Macro Conference (Southern Methodist University) and the IMIM seminar series. The usual disclaimers apply.

[†]University of New Hampshire: loris.rubini@unh.edu.

[‡]University of Cagliari and CEPR. E-mail: amoro@unica.it.

1 Introduction

In this work, we investigate the ability of workhorse structural change models in accounting for business cycle properties of an economy. Typically, structural change in general equilibrium is generated by assuming specific functional forms for preferences. Such preferences induce structural change either through a non-unitary elasticity of substitution of consumption components that drives a reallocation of resources across sectors as relative prices change, and/or through heterogeneous non-homothetic components of the various goods and services entering the utility function.[1] The most successful models in accounting for the observed sectoral reallocation include Herrendorf, Rogerson, and Valentinyi (2014) (HRV henceforth), which encompasses the specifications in Kongsamut, Rebelo, and Xie (2001) and Ngai and Pissarides (2007) as special cases; Boppart (2014) (BOP); and Comin, Lashkari, and Mestieri (2021) (CLM).[2] The key difference across these models is given by the specification of preferences. While they can all account well for the long-run properties of structural transformation in U.S. data, their ability to replicate the short-run properties along the growth path remains unexplored. This is the focus of this paper.

There a two main reasons why our exercise is of interest. First, modern business-cycle theory is grounded on the view that cycles are "oscillations of output and prices about a trend path" (Lucas, 1975), suggesting that a theory of the cycle should be embodied in a broader theory that can account for the trend (Cooley and Prescott, 1995). The one-sector exogenous growth model is succesful in providing such a theory: given permanent TFP changes, the model accounts well for the main growth facts; given stochastic and transitory TFP changes, the model accounts reasonably well for the main business-cycle facts. Structural change models are extensions of the one-sector growth model that account for additional facts of the growth process (i.e. structural transformation). Thus, a theoretical question is whether a unified theory of growth and cycles also exists for this class of models.

The second reason stems from practical considerations. There is a growing interest in the literature regarding the use of structural change models to explore phenomena over the

[1] Both channels have been shown to drive structural change in the data, with proportions that can vary depending of the preference specification. Income effects account for a proportion between 50% (Boppart, 2014) and 75% (Comin, Lashkari, and Mestieri, 2021) of structural transformation in the U.S. See Herrendorf, Rogerson, and Valentinyi (2014) and Moro and Valdes (2021) for a discussion of the two channels.

[2] More recently, Alder, Boppart, and Müller (2022) has added to this group. They characterize a flexible general class of preferences for which aggregate expenditure and saving are independent of inequality and show that these preferences account well for historical structural transformation in expediture in U.S., the U.K., Canada, and Australia. While these preferences perform very well in terms of structural transformation, we do not investigate their business-cycle properties due to the substantially larger number of parameters with respect to the other type or preferences we consider. See also Bohr, Mestieri, and Robert-Nicoud (2023), who show that a class of preferences which they label "heterothetic Cobb-Douglas" can generalize the results in Kongsamut, Rebelo, and Xie (2001).

short term. This highlights a growing recognition of the importance of considering the long-term evolution of the economy when dealing with business-cycle frequency issues. These contributions typically involve quantitative analyses incorporating various elements into the structural change block of the model to align it with the specific facts it aims to explain. However, the literature currently lacks a benchmark outcome, leaving the cyclical properties of the fundamental framework of structural transformation unestablished. Our work offers a reference point for researchers looking to apply these models in targeted studies.

Given the above considerations, our main interest is to follow as closely as possible, in accordance with the features of each model, the real business cycle (RBC) literature, by positing that the stochastic block of the economy is given by sectoral stochastic TFP processes. We thus introduce the same stochastic structure in all models, and determine the value of the relevant parameters by using data on relative prices. This allows us to compare business cycles under different preference specifications when the model economies fit the long-run structural transformation and the stochastic block is common across models.

A non-trivial complication of the analysis is that the solution of stochastic structural change models requires solving for the entire *stochastic growth path*. This happens because, under general conditions, growth in structural change models is unbalanced.[3] This implies that we cannot rely on traditional approaches that focus on a steady state or a balanced growth path (BGP) and compute deviations from the non-stochastic trend. To address this we resort to the method proposed in Fair and Taylor (1983), the Extended Path Method, which allows to solve for the whole stochastic growth path.[4]

In line with the RBC literature, we analyze both volatility and cyclicality (i.e. correlations with GDP) statistics, at the aggregate and at the sectoral level. We report four main results. First, we investigate the ability of the models to reproduce the volatility of aggregate variable observed in the data. We find that in all models this is substantially lower than in the data and in the standard one-sector RBC model. For instance, HRV accounts for 55% of GDP volatility, BOP for 45% and CLM for only 28% of it. However, all models do reasonably well

[3]The literature discusses how some structural change models may display balanced growth when measured in terms of a numeraire good, and unbalanced growth when measured in line with NIPA conventions (see Leon-Ledesma and Moro, 2020 and Duernecker, Herrendorf, and Valentinyi, 2021). Here we refer to theoretical unbalanced growth in terms of the numeraire good. See Buera, Kaboski, Mestieri, and O'Connor (2020) for an algorithm to solve deterministic structural change models outside theoretical balanced growth paths.

[4]Notable exceptions which do not resort to local methods are Cagliarini and Kulish (2013) and Maliar, Maliar, Taylor, and Tsener (2020), who develop solution methods designed for cases in which there are exogenous, deterministic structural changes, like a change in policy. In our case, structural transformation happens endogenously given the rate of growth of TFP. Notwithstanding, Storesletten, Zhao, and Zilibotti (2019) apply Cagliarini and Kulish (2013) algorithm in a stochastic environment. A necessary condition for this algorithm is that there is long run convergence to a steady state or BGP. The extended path used in our paper does not require such assumption.

in accounting for the relative volatility of consumption to GDP.

In terms of sectoral volatilities, we consider the absolute and relative performance of the models. CLM is the only model providing a good account of absolute volatilities in the data. It can explain 75% of the volatility of goods consumption and 65% of the volatility of services consumption in the data. HRV displays numbers that are substantially smaller than in the data, while BOP matches the volatility of services but displays a volatility of goods which is substantially smaller than the data. In terms of relative volatilities, we focus on a robust observation in the data which is that goods consumption is more volatile than services. Moro (2015) discusses how Stone-Geary preferences can display a tension between the ability to account simultaneously for long-run structural change and for the relative volatility of goods and services in the data. We investigate whether such tension emerges as a general feature of structural change models. Indeed, HRV displays that services consumption is extremely more volatile than goods consumption, confirming the results in Moro (2015). BOP shows a similar outcome. CLM, instead, produces a resonable 39% for the ratio of volatilities between services and goods, close to the ratio in the data, equal to 44%.

In terms of the cycle, all models do reasonably well in accounting for the cyclicality of aggregate GDP components. Both aggregate consumption and investment display a large, positive correlation with GDP. However, things are substantially different when looking at the cyclicality of sectoral variables. In terms of sectoral consumption and sectoral labor cyclicality, only HRV shows a good accounting of the data. It also performs well in accounting for the cyclicality of aggregate labor, something that the other two models cannot do. BOP accounts for only 33% of the correlation of labor and GDP, while CLM displays a small negative correlation.

Our paper relates to the growing number of contributions showing that the composition of an economy can have effects that are observable at the business cycle frequencies. Da-Rocha and Restuccia (2006) use a model with agriculture and non-agriculture sectors to show that the size of the employment share in agriculture can account for a large fraction of the differences in the magnitude of aggregate output volatility across countries. Similarly, Moro (2012) models an economy displaying structural transformation between manufacturing and services and finds that in the calibrated model the rise of the value added share of services can account for 28% of the decline in aggregate output volatility observed in the U.S. after 1980, while Moro (2015), in the context of a similar model, finds that structural transformation can account for at least 83% of the larger output volatility in middle-income relative to high-income economies. Carvalho and Gabaix (2013) perform a volatility accounting exercise for the U.S. and show that aggregate output volatility can be traced back to the change in size of the various sectors in the economy that display heterogeneous volatilities.

Yao and Zhu (2021) use a two-sector model to show that the absence of employment-output correlation in China is due to the large size of the agricultural sector in that country, while Storesletten, Zhao, and Zilibotti (2019) provide a stochastic model displaying an acceleration of structural change in booms and a deceleration in recessions. Galesi and Rachedi (2019) show that structural transformation between manufacturing and services in a New-Keynesian framework induces an increase in the share of services in intermediate goods used in the economy, which dampens the response of inflation to monetary policy shocks. Finally, Moro and Rachedi (2022) study a model where a decline of the public-sector relative productivity drives a changing structure of government spending, which modifies the transmission mechanism of government spending shocks. We add to this literature by assessing the ability of the most commonly used theories of structural change to account for business cycle properties of an economy.

The remainder of the paper is as follows. Section 2 describes the stochastic versions of the three multi-sector growth models we consider, and in section 3 we describe the solution algorithm. Section 4 describes the calibration, while section 5 presents the quantitative results. Section 6 concludes.

2 Theoretical Framework

In this section we describe the demand and supply sides of our setting. Time is discrete and there are two sectors in the economy: goods (g) and services (s). There are two representative firms, each producing one of the consumption goods. We also adopt the convention that investment is an aggregate of goods and services as in Herrendorf, Rogerson, and Valentinyi (2021) and García-Santana, Pijoan-Mas, and Villacorta (2021). For the demand side, we consider three different preference specifications: i) Herrendorf, Rogerson, and Valentinyi (2014), ii) Boppart (2014) and iii) Comin, Lashkari, and Mestieri (2021). Consistently with the RBC literature, we add leisure to these preferences. The production structure follows the standard convention of assuming a Cobb-Douglas prodution function in each sector, with the capital share identical across sectors. Once again following the RBC convention, we introduce (sector specific) TFP shocks as the only stochastic component of the environment.

2.1 Households

In this subsection we describe the three types of preferences that we analyze. For each of them, we consider an expected utility specification by introducing the expectations operator $\mathbb{E}$. There is a representative household owning the stock of capital, which is rented out in the market. The capital stock K evolves according to

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

where $\delta \in [0,1]$ is the rate of depreciation and X_t is the amount of investment chosen by the household for time t. Following Herrendorf, Rogerson, and Valentinyi (2021), the investment good is obtained by combining output of the goods and the services sectors according to the CES aggregator:

$$X_t = e^{z_{xt}} \left(\omega_x^{\frac{1}{\epsilon_x}} X_{gt}^{\frac{\epsilon_x - 1}{\epsilon_x}} + (1-\omega_x)^{\frac{1}{\epsilon_x}} X_{st}^{\frac{\epsilon_x - 1}{\epsilon_x}} \right)^{\frac{\epsilon_x}{\epsilon_x - 1}}, \tag{1}$$

where the weights ω_x and $1-\omega_x$ denote the relative importance of sector $j = g, s$ and $\epsilon_x > 0$ is the elasticity of substitution, while z_{xt} is a deterministic investment-specific TFP term.

The representative household owns an amount of labor $\bar{h}$ each period and decides how much to supply to the market in exchange for a wage w_t. Denoting with p_{jt} the price of each good for $j = g,\ s$, with p_{xt} the price of investment, and with r_t the rental rate of capital, the period t budget constraint of the household is

$$p_{gt}c_{gt} + p_{st}c_{st} + p_{xt}X_t = w_t h_t + r_t K_t, \tag{2}$$

where c_{jt} is period t consumption of good $j = g, s$, and h_t is the amount of labor supplied. We next introduce the preference specification for each model.

2.1.1 Herrendorf, Rogerson and Valentinyi (2014)

Herrendorf, Rogerson, and Valentinyi (2014) (HRV henceforth) propose a demand setting based on Stone-Geary preferences that nests Kongsamut, Rebelo, and Xie (2001) and Ngai and Pissarides (2007) as special cases. The intertemporal utility is given by

$$U = \mathbb{E} \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\tau}}{1-\tau} + \nu_t \frac{(\bar{h} - h_t)^{1-\gamma}}{1-\gamma} \right],$$

where

$$C_t = \left[\omega^{1/\mu} c_{gt}^{\frac{\mu-1}{\mu}} + (1-\omega)^{1/\mu} (c_{st} + \bar{c})^{\frac{\mu-1}{\mu}} \right]^{\frac{\mu}{\mu-1}}, \tag{3}$$

The weights ω and $1-\omega$ denote the relative importance of sector $j = g, s$, τ is the relative risk aversion, and $\mu > 0$ governs the elasticity of substitution. The term $\bar{c}$ introduces a non-homothetic component, while γ governs the Frisch elasticity. Note that, to allow labor supply to be constant along the deterministic growth path we need to allow the disutility of labor ν_t to change over time, as in Moro (2012).[5]

2.1.2 Boppart (2014)

Boppart (2014) (BOP henceforth) proposes an intratemporal indirect utility function that belongs to the "price independent generalize linearity" (PIGL) class of preferences. The intertemporal utility is given by

$$U = \mathbb{E}\sum_{t=0}^{\infty}\beta^t\left[V(p_{gt}, p_{st}, E_t) + \nu_t\frac{(\bar{h}-h_t)^{1-\gamma}}{1-\gamma}\right],$$

where

$$V(p_{gt}, p_{st}, E_t) = \frac{1}{\epsilon}\left(\frac{E_t}{p_{st}}\right)^{\epsilon} - \frac{\varphi}{\eta}\left(\frac{p_{st}}{p_{gt}}\right)^{-\eta} - \frac{1}{\epsilon} - \frac{\varphi}{\eta},$$

is an indirect utility function of c_{gt} and c_{st} with $0 \leq \epsilon \leq \eta \leq 1$ and $\varphi > 0$. E_t is consumption expenditure on goods and services. The parameter ϵ governs the evolution of the expenditure share of services, ϵ and η jointly govern the elasticity of substitution between services and goods, while φ is a shift parameter. Note that in this setting, we have a period utility given in part by an indirect utility function in expenditure and prices of goods and services, and in part by a direct utility function in labor. The dynamic problem of the household is still well specified, with the two control variables in the dynamic problem being E_t and h_t, and the budget constraint (2) is given in this case by

$$E_t + p_{xt}X_t = w_t h_t + r_t K_t.$$

As in Boppart (2014), once E_t is determined in the dynamic problem, Roy's identity can be used to find the levels of c_{gt} and c_{st}.

2.1.3 Comin, Lashkari and Mestieri (2021)

Comin, Lashkari, and Mestieri (2021) (CLM henceforth) introduce a utility function that belongs to a class of preferences defined as non-homothetic constant elasticity of substitution (CES) preferences. The intertemporal utility is

$$U = \mathbb{E}\sum_{t=0}^{\infty}\beta^t\left[\frac{C_t^{1-\tau}}{1-\tau} + \nu_t\frac{(\bar{h}-h_t)^{1-\gamma}}{1-\gamma}\right],$$

[5] We adopt the same convention for all models.

as in HRV. However, here aggregate consumption C_t is defined implicitly as follows

$$\left[\psi_g^{1/\sigma} C_t^{\frac{\epsilon_g - 1}{\sigma}} c_{gt}^{\frac{\sigma-1}{\sigma}} + \psi_s^{1/\sigma} C_t^{\frac{\epsilon_s - 1}{\sigma}} c_{st}^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} = C_t,$$

where we are using the specification in Duernecker, Herrendorf, and Valentinyi (2023) and Duarte and Restuccia (2023). The weights ψ_g and ψ_s denote the relative importance of sector $j = g, s$ in consumption, ϵ_g and ϵ_s govern the income elasticities of the two types of consumption, and $\sigma \geq 0$ controls the elasticity of substitution between goods and services.

2.2 Firms and Market Clearing

Each good j is produced in sector j with a technology that employs as inputs capital and labor. Such technology is given by

$$Y_{jt} = e^{z_{jt}} K_{jt}^{\alpha} L_{jt}^{1-\alpha},\ 0 \leq \alpha \leq 1,\ j = g, s,$$

where K_{jt} and L_{jt} are capital and labor used in sector j and z_{jt} is a sector-specific stochastic TFP term. We retain the convention of several structural change models in assuming that the capital and labor shares are identical in all sectors of the economy.

Feasibility requires the following market clearing conditions at any t:

$$Y_{gt} = c_{gt} + X_{gt}, \quad Y_{st} = c_{st} + X_{st}, \quad K_t = \sum_{j=g,s} K_{jt}, \quad h_t = \sum_{j=g,s} L_{jt}.$$

2.3 Stochastic Processes

The evolution of technology in each sector j, z_{jt}, is given by the sum of a first order autoregressive process, plus a deterministic component that grows with time. More specifically,

$$\begin{aligned} z_{gt} &= g_g t + \hat{z}_{gt}, \quad & \hat{z}_{gt} &= \rho_g \hat{z}_{gt-1} + \epsilon_{gt}, \quad & \epsilon_{gt} &\sim N(0, \sigma_g^2) \\ z_{st} &= g_s t + \hat{z}_{st}, \quad & \hat{z}_{st} &= \rho_s \hat{z}_{st-1} + \epsilon_{st}, \quad & \epsilon_{st} &\sim N(0, \sigma_s^2) \end{aligned}$$

where g_j is a deterministic trend, ρ_j is the autoregressive parameter, and σ_j^2 is the variance of shocks ϵ_{jt} for each sector $j = g, s$. Perfect competition implies that relative prices are uniquely determined by TFP. More precisely, taking the price of goods as the numeraire at each t, we have

$$p_{gt} = 1, \quad p_{st} = e^{z_{gt} - z_{st}}.$$

These equations allow us to estimate the parameters governing the stochastic component of TFP independently of the choice of preferences, making our results comparable across models.

3 Solving the models

There is no standard procedure for solving stochastic structural change models along the transition path. Under general conditions, structural change implies an unbalanced growth path, and this property prevents the use of local solution techniques when uncertainty is introduced, requiring the adoption of global methods.[6] We adopt the methodology proposed by Fair and Taylor (1983), the Extended Path Method. This method consists in first finding a deterministic dynamic equilibrium of the model over the period of interest. In practice, if the last period of interest is denoted T_0, the algorithm extends the period of analysis by ΔT periods, so that the deterministic dynamic equilibrium is found until period $T_1 = T_0 + \Delta T$. The second part of the algorithm computes the stochastic path given the expectation of the different shocks. In particular, the assumption is that, for any shock occurred up to time T_0, the economy rests in the deterministic equilibrium in period T_1. Thus, the size of ΔT must be large enough so that this does not affect the equilibrium in the interval between 0 and T_0. We provide more details of the methodology next.

3.1 Deterministic Path

To find the deterministic path, we use a "shooting algorithm".[7] Similarly to a one-sector growth model, the solution implies solving a second order difference equation subject to two boundary conditions: an initial condition for the capital stock, and a transversality condition holding at $t \to \infty$. Typically, in a standard one-sector growth model, the transversality condition is replaced with a steady state level for capital as $t \to \infty$. However, our models do not necessarily have a steady state, so we need to impose the transversality condition as a boundary condition. A complication is then that this condition should be satisfied only as

[6] Among the preferences specifications we consider, only Boppart (2014) allows the existence of a balanced growth path with a finite capital stock and constant sectoral TFP growth. The preferences in Herrendorf, Rogerson, and Valentinyi (2014) admit a balanced growth path with finite capital stock and constant sectoral TFP growth only in the special cases described in Kongsamut, Rebelo, and Xie (2001) and Ngai and Pissarides (2007).

[7] See Ljungqvist and Sargent (2004), chapter 11, for details of the algorithm.

$t \to \infty$, which cannot be computed numerically. Consequently, we approximate ∞ with T_1 and evaluate the transversality condition in that period.

3.2 Stochastic Path

After successfully finding the deterministic path, we compute the stochastic path. The key assumption the Extended Path method relies on is approximating the expectation of a function with a function of the expectations.[8] In practice, we use the shooting algorithm repeatedly, each time under different (known) sequences of shocks, with the constraint that the equilibrium allocation should coincide with the deterministic path in period T_1.

In practice, the method works as follows. Starting in period 0, we observe the realization of the shocks (i.e. innovations ϵ_{g0} and ϵ_{s0}). This allows us to compute the expectation of $\hat{z}_{jt}$'s in all future periods $t = 1, \ldots, T_1$. As innovations have zero mean, the expectation at time 0 is that the $\hat{z}_{jt}$'s will mean revert to zero following their autoregressive processes over time. These processes are then treated as deterministic sequences from the standpoint of time 0. It follows that we can use the shooting algorithm to compute the equilibrium allocation until T_1 given k_0 and this sequence of $\hat{z}_{jt}$'s. From this dynamic equilibrium we save decision rules for consumption and investment in period 0. Next, we move on to period 1. Notice that k_1 is known at this point. We obtain new realizations for the stochastic processes (i.e. innovations ϵ_{g1} and ϵ_{s1}). Given these realizations, we compute the expected $\hat{z}_{jt}$'s for $t = 2, \ldots, T_1$. Given these, we again use the shooting to compute the equilibrium allocations assuming the $\hat{z}_{jt}$'s will deterministically go back to zero over time and save the decision rules for consumption and investment in period 1.

We continue until period T_0 is reached. In this way, the decision rule for period $t \leq T_0$ is the equilibrium allocation obtained in the step that starts in period t. This illustrates the importance of choosing a large enough ΔT. A small value would bias the results, since it would affect the decision rules in periods $t \leq T_0$. A large ΔT avoids this problem.[9]

4 Parameter Values and Measurement

This section describes how we assign values to the parameters in the model and how we compute aggregate variables in the models in a way that allows comparison with the data.

[8] Note that this approximation is perfect only when the functions are linear, which is typically not the case. However, this assumption is similar to what most RBC models employ to approximate the rational expectation equilibrium, using first order Taylor expansions.

[9] We follow the methodology in Fair and Taylor (1983) to select ΔT: this is chosen such that adding one more period would not change the equilibrium until period T_0 in a substantial way.

Given that our theoretical setup is similar, we closely follow the two-by-two data treatment in Herrendorf, Rogerson, and Valentinyi (2021), who combine U.S. industry data from WORLD KLEMS (value added, factor shares and factor inputs) with input–output tables from the Bureau of Economic Analysis.

Pre-determined parameters: One period in the model represents one year in the data. Consequently, we use standard values for the following parameters when one period is one year: the capital depreciation rate is $\delta = 0.08$, and the subjective discount factor is $\beta = 0.96$. In addition, following standard practices we set the capital share $\alpha = 1/3$ and the value of the relative risk aversion $\tau = 3$ in HRV and CLM.[10]

Preference parameters for consumption: We estimate preference parameters as in Herrendorf, Rogerson, and Valentinyi (2013) by using iterated feasible generalized nonlinear least square estimation.[11] The equations to estimate are all derived in the papers that introduce those preferences. For convenience, we re-write them here. Let η_{jt} denote the share of expenditure on sector $j = g, s$ at time t. This is observable in the data, along with relative prices and expenditures. In all models, these shares are functions of the parameters, prices, and expenditures. Note, however, that shares add up to one, so that we cannot estimate both equations due to linear dependency, and we perform the estimation of only one equation.[12]

The following equation corresponds to equation (14) in Herrendorf, Rogerson, and Valentinyi (2013) for our case with two sectors:

$$\eta_{st} = \frac{\frac{1-\omega}{\omega} p_{st}^{1-\mu} - p_{st}\bar{c}/E_t}{1 + \frac{1-\omega}{\omega} p_{st}^{1-\mu}}.$$

For BOP we estimate

$$\eta_{st} = 1 - \varphi E_t^{-\epsilon} p_{st}^{\epsilon-\eta}.$$

In the case of CLM, we are adopting the specification in Duernecker, Herrendorf, and Valentinyi (2023) and Duarte and Restuccia (2023). Appendix B derives the following expression for estimation:

$$\log\left(\frac{\eta_s}{\eta_g}\right) = \log\left(\frac{\Psi_s}{\Psi_g}\right) + \frac{\epsilon_s - \epsilon_g}{\epsilon_g - \sigma}\left[\log(\eta_g) - (\sigma - 1)\log(E) - \log \Psi_g\right] + (1-\sigma)\log(p_s). \quad (4)$$

[10] A relatively large value of τ guarantees finding a solution in CLM given our computational algorithm. See Appendix A for details. We then use the same value for HRV.

[11] Deaton (1986).

[12] The choice of the equation is irrelevant for the estimation purposes.

As in Comin, Lashkari, and Mestieri (2021) we normalize $\epsilon_g = \Psi_g = 1$ to obtain

$$\log\left(\frac{\omega_s}{\omega_g}\right) = \log\left(\Psi_s\right) + \frac{\epsilon_s - 1}{1-\sigma}\log(\omega_g) + (\epsilon_s - 1)\log(E) + (1-\sigma)\log(p_s).$$

We estimate equation 4 by imposing $\epsilon_s = \sigma$.[13]

The estimated parameter values are reported in the top panel of Table (1). In what follows we compare the values we obtain with other calibrations and estimations in the literature in a similar context.

Our estimates for HRV are $\mu = 0.19$, $\omega = 0.13$ and $\bar{c} = 1.80$. Similar preferences are calibrated in Moro (2012) for the U.S. for the 1960-2005 period. He obtains 1.27 for the non-homotetic term $\bar{c}$, and a small value (0.000023) for the weight of goods ω. Also, he follows Rogerson (2008) and Duarte and Restuccia (2010) in imposing a value of μ of 0.4 which implies that consumption of goods and consumption of services are complement, as in our estimate. It is worth highlighting that finding $\bar{c} > 0$ implies that there is a subsistence level of consumption in goods, so that, given a low level of income, a large proportion of it will be spent on goods. As income increases, the proportion spent on goods loses ground to that of services.

The values we obtain for BOP are $\epsilon = 0.56$, $\eta = 0.67$ and $\varphi = 0.37$. Boppart (2014) provides estimates obtained with micro data of the the first two parameters, given by 0.22 for ϵ and 0.41 for η. Leon-Ledesma and Moro (2020) instead use macro data to calibrate BOP for the U.S. using the 1951-2015 period and obtain $\epsilon = 0.22$, $\eta = 0.51$ and $\varphi = 0.64$. All exercises report values that are fairly close to each other. To have a better understanding of their meaning, it is convenient to write down the demand for goods implied by these preferences:

$$c_{gt} = \varphi E_t^{1-\epsilon} p_{st}^{\epsilon-\eta}$$

Our estimates for ϵ are somewhat larger than the other papers mentioned, implying a stronger income effect in our model than in the other works. The difference $\epsilon - \eta$ is -0.11 in our case, similar to that in Boppart (2014) and somewhat smaller in absolute value (but of the same sign) than Leon-Ledesma and Moro (2020).

Finally, we obtain the following estimates for CLM: $\sigma = 1.27$, $\epsilon_s = 1.27$, $\psi_s = 0.66$, plus the normalized $\epsilon_g = \psi_g = 1$. A recent paper by Duarte and Restuccia (2023) calibrates a

[13] Without imposing $\epsilon_s = \sigma$, we obtain estimates $\sigma = 1.11$ and $\epsilon_s = 1.12$. While the estimated values of the two parameters are close, the shooting algorithm used to solve the model with the extended path methodology fails due to non-monotonicities in some components of the Euler equation. For these reasons, we impose $\epsilon_s = \sigma$ in the estimation procedure, obtaining a value of 1.27 for both parameters. While this is slighly larger than the values of the two parameters obtained in the unconstrained estimation, the value of Ψ_s is also slightly different with respect to that case. See Appendix A for details.

two-sector model in goods and services with CLM preferences for the U.S. for the period 1960-2000. They calibrate the difference in income effects to $\epsilon_s - \epsilon_g = 1$, while we obtain a difference of 0.27. In addition, while they do not normalize $\psi_g = 1$, they calibrate the ratio $\psi_s/\psi_g = 0.69$, while we have $0.66/1 = 0.66$. Finally, they impose a value of $\sigma = 0.1$, while our estimate gives a substantially larger value of 1.27. Thus, our estimation of CLM implies that goods and services are relative substitutes in the utility function, something in contrast with previous estimates in Comin, Lashkari, and Mestieri (2021) and Duernecker, Herrendorf, and Valentinyi (2023) in the context of three-sector models.[14]

Table 1: Parameter values across models

HRV			BOP			CLM		
	Value	**S.E.**		**Value**	**S.E.**		**Value**	**S.E.**
μ	0.1931	(0.0929)	ϵ	0.5567	(0.0375)	σ	1.2702	(0.0076)
ω	0.1326	(0.0134)	η	0.6698	(0.0715)	ϵ_s	1.2702	(0.0076)
$\bar{c}$	1.8012	(0.2884)	φ	0.3673	(0.0029)	ψ_s	0.6608	(0.0038)
						ϵ_g	1	(-)
						ψ_g	1	(-)
	Value	**Target**		**Value**	**Target**		**Value**	**Target**
γ	0.21	Volat. of L	γ	0.12	Volat. of L	γ	0.24	Volat. of L
ν_0	0.0467	Const. L	ν_0	1.0211	Const. L	ν_0	0.0016	Const. L
ν_{T_1}	0.0183		ν_{T_1}	2.0604		ν_{T_1}	$1.3E-07$	
	Value	**Target**		**Value**	**Target**		**Value**	**Target**
g_g	0.0148	η_{s1947}	g_g	0.0145	η_{s1947}	g_g	0.0121	η_{s1947}
g_s	0.0070	g_{p_s}	g_s	0.0068	g_{p_s}	g_s	0.0044	g_{p_s}
k_0	2.3810	η_{s2014}	k_0	2.7992	η_{s2014}	k_0	0.5011	η_{s2014}

Note: η_{s1947} and η_{s2014} are the shares of services in consumption in the data in the two years. g_{p_s} is the average growth rate of the relative price of services to goods in the data in the sample period 1947-2014.

Preference parameters for labor/leisure: We calibrate γ such that each model matches an average volatility of aggregate labor as in the data, which is 2.22%. For ν_t we impose a time sequence ensuring that, along the deterministic growth path, the supply of labor in the market is 1/3 of total time of the household at each time t. We set the total labor endowment each period $\bar{h} = 3$, so that in the deterministic equilibrium labor supply equals

[14] The fact that the estimate for σ is larger than one does not depend on imposing the constraint $\sigma = \epsilon_s$ or the normalization $\epsilon_g = \Psi_g = 1$.

1 each period. This implies a declining value for ν_t in HRV and CLM and an increasing one in BOP. The middle section of Table (1) reports the initial and the final value of ν_t for each model.

Initial capital and TFP growth: Assume we know the values of k_0, g_g, and g_s. Given the pre-determined parameters and the preference parameters, we can solve for the deterministic equilibrium path that satisfies k_0 and the transversality condition. To pin down values for k_0, g_g, and g_s we use the following targets: i) the share of services in consumption in the initial year (1947), ii) in the final year (2014), and iii) the growth in the relative price of services between 1947 and 2014, as in equilibrium $p_{st} = e^{z_{gt}-z_{st}}$. The bottom section of Table (1) reports these parameter values for each model.

Parameters of the investment function: From Herrendorf, Rogerson, and Valentinyi (2021), we observe that the price of investment relative to goods does not display a clear trend over time.[15] We thus compute the average value over the sample, obtaining 1.13, and impose that the model matches that target in each period of the deterministic transition (i.e. that there is no trend in the relative price of investment to goods along the growth path, as in the data). We then note that in all models the price can be written as

$$P_{xt} = \left[\frac{\omega_x + (1-\omega_x)P_{st}^{1-\epsilon_x}}{e^{z_{xt}(1-\epsilon_x)}}\right]^{1/(1-\epsilon_x)}.$$

Imposing that $P_{xt} = 1.13$ at each t, we can use the above formula together with the time series of the relative price of services P_{st} to back out the sequence of z_{xt}. In Table 2 we report the average value of z_{xt}, which implies a value of the "productivity" term in the investment function $e^{z_{xt}} = 0.98$.

Next, from the first order conditions of the cost minization problem of investment we can write, after taking logs, the following expression

$$\log\frac{p_{st}X_{st}}{X_{gt}} = (1-\epsilon_x)\log p_{st} + \log\frac{1-\omega_x}{\omega_x}. \tag{5}$$

Equation (5) can be estimated via OLS using the investment share of services relative to goods and the relative price of services to goods to obtain values for ϵ_x and ω_x.[16] Note that all parameters of the investment function are common across models. Herrendorf, Rogerson, and Valentinyi (2021) estimate equation (5) obtaining the values of ϵ_x and ω_x that we report in Table 2, which we use in our simulations.

[15] See Figure 2 in Herrendorf, Rogerson, and Valentinyi (2021).

[16] The composition of investment is decided by the household in her maximization problem. However, the resulting input composition is equivalent to that obtained by minimizing the cost of producing one unit of investment, which we exploit here.

Table 2: Parameter values for Investment

Parameter	Value	Source
$\bar{z}_x$	0.0136	Target p_x
ϵ_x	0.0002	OLS
ω_x	0.6482	OLS

Note: $\bar{z}_x$ is the average value along the transition path of z_{xt}.

TFP Processes: To calibrate the parameters $\rho_g, \rho_s, \sigma_g, \sigma_s$ we estimate AR(1) processes using data on real sectoral output, sectoral employment and sectoral capital. We proceed as follows:

1. Compute the growth rate of sectoral TFP. We do this by log-linearizing production functions as follows:

$$g_{jt} = \log(z_{jt} - z_{jt-1}) = \log\left[\frac{Y_{jt}}{Y_{jt-1}}\right] - \alpha \log\left[\frac{K_{jt}}{K_{jt-1}}\right] - (1-\alpha)\log\left[\frac{L_{jt}}{L_{jt-1}}\right], j = g, s.$$

2. Detrend the series for these growth rates with a linear trend. We use a linear trend because the theoretical trend is linear in logs. The linear trend is computed by estimating the following regression using OLS:

$$g_{jt} = a_0 + a_1 t$$

3. Compute the residuals as

$$r_{jt} = g_{jt} - \hat{a}_0 - \hat{a}_1 t$$

   where $\hat{a}_0$ and $\hat{a}_1$ are the estimates of a_0 and a_1, respectively.

4. Estimate the following autorregressive process:

$$r_{jt+1} = \rho_j r_{jt} + \epsilon_{jt}, \quad j = g, s, \quad \epsilon \sim N(0, \sigma_j^2).$$

Table 3 reports parameter values. The estimated autoregressive parameters are $\rho_g = 0.91$ and $\rho_s = 0.96$. These are in line with estimates in Moro (2012), in the context of a similar two-sector split of the economy. The volatility estimates imply that the stochastic component of goods TFP is more volatile than its services counterpart. Conditional on current TFP, the volatility of next period's TFP is twice as large in goods than in services (0.0227/0.0111). The unconditional variance in goods ($\sigma_g^2/(1-\rho_g^2) = 0.0029$) is also larger than in services ($\sigma_s^2/(1-\rho_s^2) =$

Table 3: Estimated TFP processes

Parameter	Estimated Value	S.E.
ρ_g	0.9063	(0.0472)
ρ_s	0.9584	(0.0378)
σ_g	0.0227	(0.0021)
σ_s	0.0111	(0.0010)

0.0015). This is consistent with goods being more volatile than services in the data, and again consistent with the estimated volatility of sectoral shocks in Moro (2012).[17]

Measurement: An issue that arises in multi-sector growth models is how to compute aggregate variables so that they can be compared with the data. In some models, aggregate consumption is explicitly defined (as in HRV), in some others it is implicitly definced (as in CLM) while in others it is not defined at all (BOP). In the data, aggregate consumption and GDP are measured using chain-weighted Fisher indices. This index can also be constructed from the model, using equilibrium quantities and prices of the various sectors. The literature suggests this is the best way to compare aggregate outcomes in multi-sector growth models to the data.[18] We thus follow this methodology for each model. To construct real GDP, we use a chain-weighted Fisher quantity index of goods and services value added (i.e. output of the two sectors). To construct real aggregate consumption we use a chain-weighted Fisher quantity index of goods consumption and services consumption value added (i.e. output of the two sectors). The investment index is defined in the same fashion within each model by equation (1). We use its equilibrium values to compute business cycle statistics for investment in the model.

Figure 1 reports the fit in terms of GDP, consumption and investment in the average of 1,000 simulations. All models account reasonably well for the evolution of the three macro-variables along the growth path.

[17] We also estimated AR(1) processes allowing for a cross-correlation of shocks. When feeding the models with these processes instead of the ones described in the main text, results do not change considerably with respect to those presented in Section 5. This is due to the low estimated correlation of innovations. This confirms the finding in Carvalho and Gabaix (2013) that the average correlation in TFP innovations across different sectors in the U.S. is very small and likely to be due to measurement error and factor hoarding.

[18] See Leon-Ledesma and Moro (2020) and Duernecker, Herrendorf, and Valentinyi (2021).

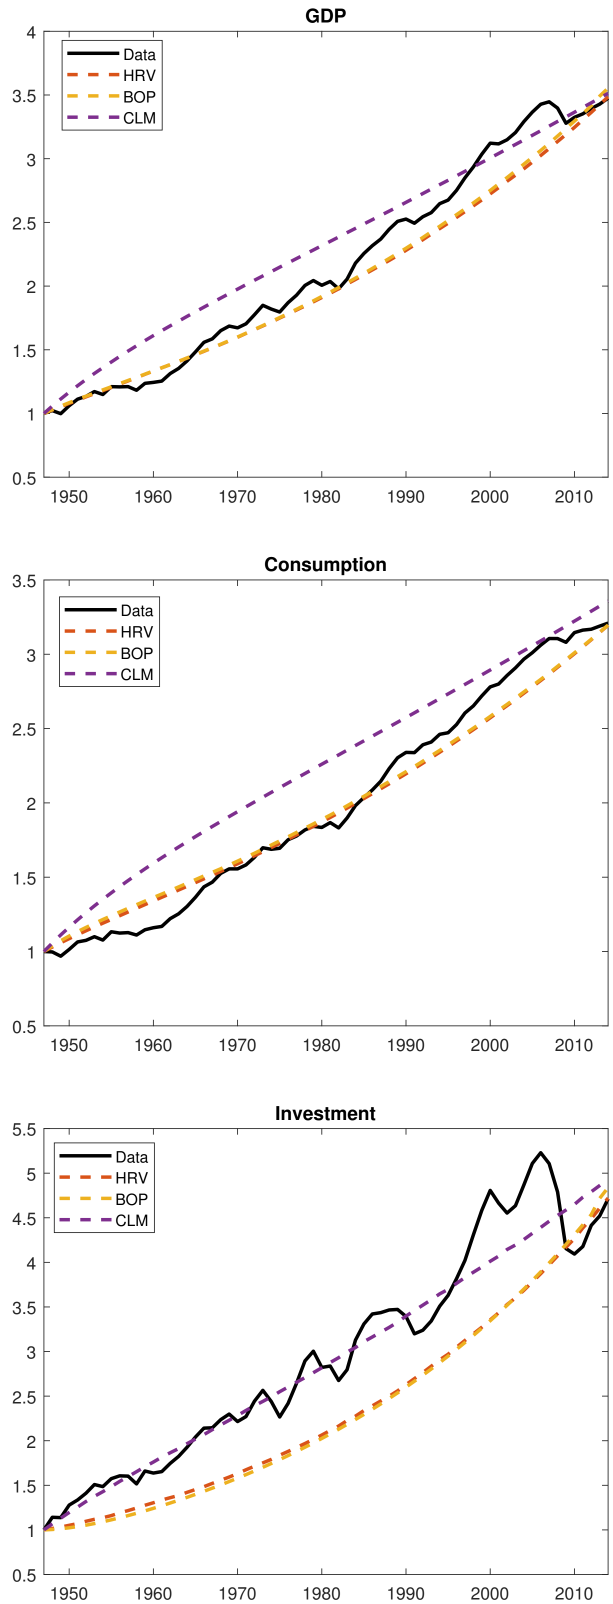

Figure 1: GDP, Consumption and Investment along the growth path. Averages across 1,000 simulations.

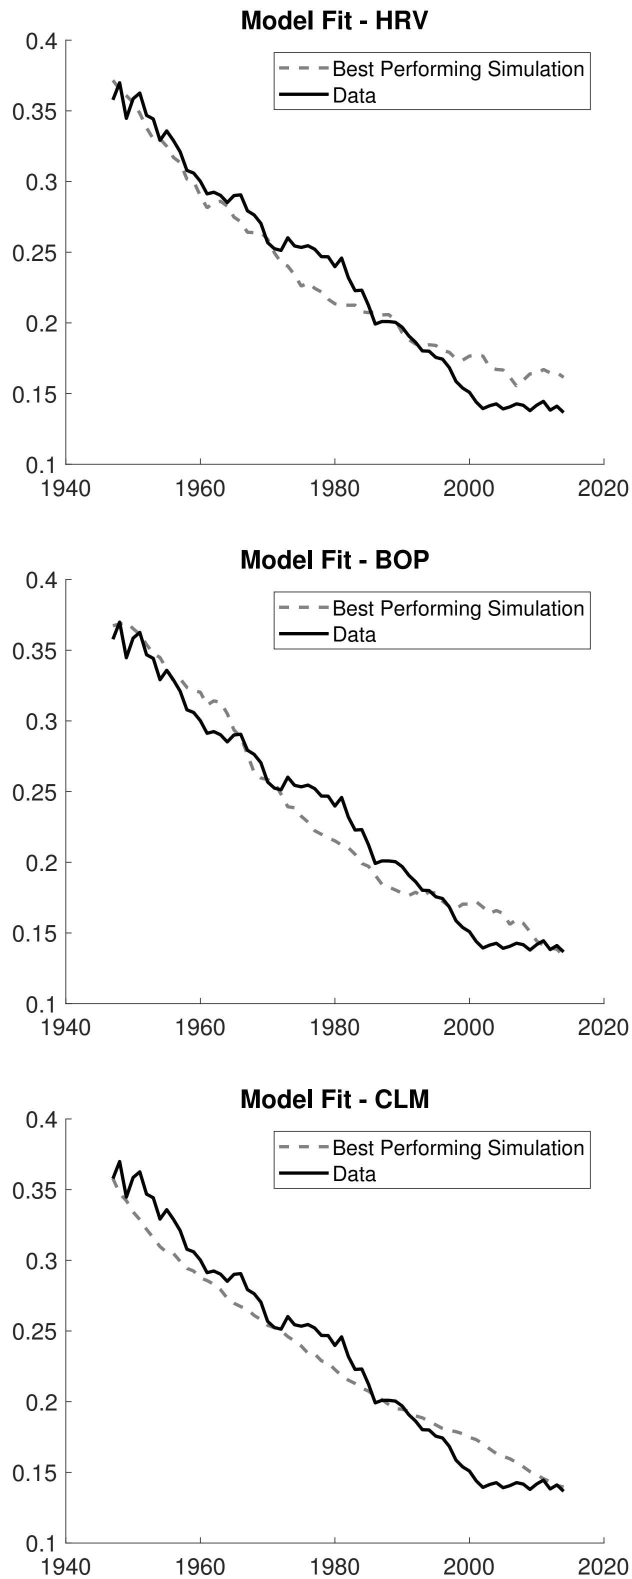

Figure 2: Structural change in a stochastic environment. Top: HRV; middle: BOP; bottom: CLM. The best performing simulation among all simulations is given by the one that minimizes the sum of square differences between model and data in the share of goods along the growth path.

5 Results

For each model, we perform 1,000 simulations that generate a sequence of TFP shocks in each sector and then compute average business cycle statistics across simulations. We first assess the ability of the models in accounting for structural change in the stochastic environment. Figure 2 shows the evolution of the value added share of goods for one of the 1,000 simulations (dashed lines), together with the corresponding figures in the data (solid lines). For each model, all simulations perform similarly in terms of long-run structural change.[19] The good performance of the models in replicating long-run-structural change is well know from the original works of Herrendorf, Rogerson, and Valentinyi (2013), Boppart (2014) and Comin, Lashkari, and Mestieri (2021). Here we confirm that a *stochastic version* of the model also fits structural change data closely. We next present the results, organized into two subsections. The first studies the volatilities generated by the models. The second studies cyclicality, that is, the co-movement of the different series with GDP.

5.1 Volatility

We start by investigating volatility statistics, first focusing on aggregate variables, whose performance is directly comparable to that in a one-sector business cycle model, and then turn to discuss sectoral statistics. Table 4 reports standard deviations for the residual component in the three specifications of the stochastic environment that we analyze and for the data.[20] In terms of absolute volatilities, all models display a substantially lower volatility of consumption, investment and GDP with respect to the data. Results for GDP, for instance, go from 55% of the volatility in the data for HRV, to only 28% in CLM, with BOP in the middle. For comparison, the standard one-sector growth model accounts for roughly 78% of U.S. GDP volatility in Cooley and Prescott (1995).

In terms of relative volatilities, a well known feature of business-cycles statistics is that investment is more volatile than GDP in the U.S., which in turn is more volatile than consumption. All three models are consistent with investment being the most volatile component. HRV and CLM also display a volatility of GDP larger than consumption, while BOP fails in this dimension. Regarding specific ratios, in the data consumption displays 76% of the volatility of GDP. In HRV consumption shows 58% of GDP volatility and in CLM it shows 69% of it. Thus, both models account for a substantial fraction of relative

[19] In Figure 2 we report the simulation which is closest to the data when using the sum of square differences as a measure of distance. However, most simulations display a similar pattern.

[20] We construct statistics as follows. First, we take logarithms of the data and detrend using a Hodrick-Prescott (HP) filter with smoothing parameter 100 (annual data). For the series produced by the model we also take logs and detrend using the HP filter with smoothing parameter 100.

consumption volatility. BOP, on the other hand, displays a volatility of consumption which is similar to that of GDP. Investment instead, is 2.36 times more volatile than GDP in the data. The best performance here is given by BOP, which displays a 1.84 larger volatility of investment relative to GDP. Both HRV and CLM overshoot substantially this statistic (3.93 and 6.12, respectively).

Consider next sectoral consumption volatility. In terms of absolute volatilities, HRV displays numbers that are substantially smaller than in the data, while BOP matches the volatility of services but displays a volatility of goods one order of magnitude smaller than the data. CLM instead, does a resonable job in accounting together for 75% of the volatility of c_m and 65% of the volatility of c_s in the data. Thus, in terms of absolute sectoral consumption volatilities this is the only model providing a reasonable account of the data.

Perhaps more important than absolute volatilities is the ability of the models to account for relative volatilities across sectors. A robust observation in the data is that goods consumption is more volatile than services.[21] However, Moro (2015) discusses how Stone-Geary preferences, like the ones in HRV, display a tension between the ability of the model to account simultaneously for long-run structural change and for the relative volatility of goods and services in the data. Table 4 allows us to investigate whether such tension between the long- and short-run properties emerges also in other structural change models. In the data, services consumption is only 44% as volatile as goods consumption. In HRV, services consumption is considerably more volatile than goods consumption, confirming the results in Moro (2015). BOP shows a similar outcome. CLM, instead, produces a resonable 39% for the ratio of volatilities between services and goods. Thus, in terms of both absolute and relative volatility of sectoral consumpion, CLM is the only specification that accounts for a substantial fraction of data statistics.

To conclude this section, consider labor volatility. Recall that the volatility of aggregate labor is matched by construction in all models, as it is used as a target in the calibration. Regarding sectoral labor volatility, this is 0.0363 for goods and 0.0249 for services in the data. All models display a substantially smaller volatility of both, except BOP, which shows a similar volatility of services as in the data. Regarding relative volatility, this is similar to that of sectoral consumpion: in the data goods labor volatility is 46% larger than that of services. Only CLM comes close to this with a 29% larger goods volatility (0.0135 versus 0.0105), while HRV displays similar numbers across sectors (0.0086 versus 0.0079) and BOP shows a volatility of services which is 33% larger than that of manufacturing (0.0276 versus

[21]This data regularity induced researchers to investigate whether the so called “great moderation” in the U.S. could be attributed to a structural change between goods and services. See McConnell and Perez-Quiros (2000), Alcalá and Sancho (2004) and Moro (2012).

Table 4: Standard deviations

	p_{st}	c_{gt}	c_{st}	C_t	X_t	Y_t	h_t	L_{gt}	L_{st}
Data	0.0249	0.0308	0.0136	0.0168	0.0520	0.0218	0.0222	0.0363	0.0249
HRV	0.0230	0.0041	0.0088	0.0069	0.0472	0.0120	0.0217	0.0086	0.0079
	(0.0028)	(0.0006)	(0.0013)	(0.0010)	(0.0059)	(0.0015)	(0.0597)	(0.0012)	(0.0010)
BOP	0.0230	0.0033	0.0138	0.0107	0.0182	0.0099	0.0221	0.0207	0.0276
	(0.0028)	(0.0004)	(0.0016)	(0.0012)	(0.0141)	(0.0027)	(0.0235)	(0.0051)	(0.0040)
CLM	0.0230	0.0231	0.0089	0.0043	0.0380	0.0062	0.0229	0.0135	0.0105
	(0.0028)	(0.0028)	(0.0011)	(0.0004)	(0.0046)	(0.0006)	(0.0292)	(0.0018)	(0.0014)

Note: The numbers reported in the table are computed as the standard deviation of the log-deviations of each variable from its Hodrick-Prescott filtered series . Standard deviations across simulations are in parenthesis. Real GDP (Y) is computed using a chain-weighted Fisher quantity index of sectoral value added. Aggregate consumption (C) is computed using a chain-weighted Fisher quantity index of sectoral consumptions.

0.0207).

5.2 Cyclicality

We now turn to study the cyclicality of endogenous variables. We stick to following the convention of the RBC literature and investigate the correlation of percentage deviations from trend of all variables with those of real GDP. These are reported in Table 5, while in Appendix C we report all cross-correlations among endogenous variables.

In term of the cyclicality of aggregate consumption and investment, all models display a positive correlation with GDP, as in the data. The positive correlation in the data is high (0.90 for consumption and 0.86 for investment). HRV matches the correlation of investment but accounts for around 63% of the correlation of consumption in the data. BOP instead, matches almost all of the correlation of consumption, but acconts for only 48% of the correlation of investment. CLM accounts for 76% of the correlation of consumption and 92% of that of investment. Thus, all models perform reasonably well along this dimension.

In terms of sectoral consumption, both components display a high correlation with GDP, 0.73 for goods and 0.82 for services. HRV shows a correlation for goods of 0.90, thus overshooting the data statistic, and of 0.47 for services, thus accounting for 57% of the data. BOP and CLM do substantially worse, as they can account for only one of the two components. In the former, the correlation of services with GDP is almost perfectly matched, but the model can account for only 21% of the correlation of goods. For CLM instead, the correlation of goods is matched reasonably well (86% accounted for) while there is basically zero correlation of services with GDP.

Table 5: Cross-Correlation with GDP

	c_{gt}	c_{st}	C_t	X_t	h_t	L_{gt}	L_{st}
Data	0.7295	0.8227	0.9040	0.8619	0.8757	0.8542	0.7971
HRV	0.9041	0.4658	0.5667	0.8877	0.6516	0.4412	0.5778
	(0.0286)	(0.1248)	(0.1052)	(0.0329)	(0.0819)	(0.0127)	(0.0109)
BOP	0.1510	0.8412	0.8467	0.4134	0.2886	0.7654	-0.0011
	(0.1638)	(0.1444)	(0.1439)	(0.3168)	(0.1658)	(0.1003)	(0.1950)
CLM	0.6309	0.0138	0.6836	0.7852	-0.2080	0.0482	-0.2081
	(0.0910)	(0.1539)	(0.0721)	(0.0435)	(0.1483)	(0.1552)	(0.1413)

Note: The numbers reported in the table are computed as the standard deviation of the log-deviations of each variable from its Hodrick-Prescott filtered series. Standard deviations across simulations are in parenthesis.

In terms of the cyclicality of aggregate labor, only HRV accounts for a reasonable fraction of the data (74%). BOP accounts for only 33% of the correlation of labor and GDP, while CLM displays a negative correlation of -0.21. The low and the negative correlation of labor with GDP in BOP and CLM, respectively, helps explaining the low volatility of GDP in Table 4 for the two models. In fact, given the common production structure of all models, one can define an aggregate production function of the economy mapping GDP to total capital and labor used in production. Once this is defined, the volatility of total labor becomes part of the volatility of GDP. If labor is positively correlated with GDP, its volatility would increase GDP volatility, if it is negatively correlated, its volatility would dampen it.

HRV appears the best model also in terms of sectoral labor cyclicality, accounting for 52% of the correlation of goods labor, and 73% of the correlation of services labor with GDP. BOP accounts for most of the correlation of goods labor, but shows no correlation of services labor with GDP. CLM performs the worst with a correlations which is roughly zero for goods, and negative for services.

6 Conclusions

The literature on structural transformation grew dramatically in the last two decades, with new models that can capture well the long run shift of resources across sectors. However, regardless of the growing use of these models in applications dealing with phenomena in the short-term, their ability to account for basic business-cycle facts has been unexplored until now. This paper makes a first step in this direction, comparing three workhorse models of structural transformation among them and relating to the RBC literature.

One result is that none of these models, in their simplest version, provides a unified theory of growth and cycles that can be regarded as a multi-sector analogue of the one-sector growth model. The latter is succesful in providing such a theory: given permanent TFP changes, the model accounts well for the main growth facts; given stochastic and transitory TFP changes the model accounts reasonably well for the main business-cycle facts. Structural change models are extensions of the one-sector growth model that account for additional facts of the growth process. This implies that they are also required to match a larger set of statistics when confronted at the business cycle frequency. We show that, probably not so unexpectedly, their ability to match these statistics is limited. Yet, our results can be used to pick the most appropriate structural change model for the specific type of short-run analysis needed. For instance, if a researcher is mainly interested in sectoral volatility, the specification that provides the best account of the data is CLM. Instead, if one is mainly interested in a setting able to reproduce the cyclicality of most components of GDP, then HRV appears as the best choice.

We stress here that richer versions of these model can potentially improve their ability to fit business-cycle facts. For instance Alder, Boppart, and Müller (2022) propose a class of intertemporally aggregable (IA) preferences that nest HRV and BOP preferences as special cases. These preferences become especially relevant to fit the data when the manufacturing share displays an inverted U-shape over time. However, they display a substantially larger number of parameters with respect to the other type or preferences we consider, which provides more flexibility, but also induce a more complicated setting. Also, in CLM the price elasticity of demand is the same for goods and services. This feature might contribute to the poor performance of the model in some dimensions, and can potentially be amended by resorting to a more general version of the CLM-preference structure which allows the price elasticity of demand to be different across goods (Hanoch, 1971, Hanoch, 1975) or across income levels (Sato, 1975). At the same time, the introduction of other type of shocks, like demand shocks, can also potentially improve the business-cycle performance of the basic models. We leave this exploration for future research.

References

Alcalá, F., and I. Sancho (2004): “Output composition and the US output volatility decline,” *Economics Letters*, 82(1), 115–120.

Alder, S., T. Boppart, and A. Müller (2022): “A theory of structural change that can fit the data,” *American Economic Journal: Macroeconomics*, 14(2), 160–206.

Bohr, C., M. Mestieri, and F. Robert-Nicoud (2023): *Heterothetic Cobb Douglas: Theory and Applications*. Centre for Economic Policy Research.

Boppart, T. (2014): “Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences,” *Econometrica*, 82(6), 2167–2196.

Buera, F. J., J. P. Kaboski, M. Mestieri, and D. G. O’Connor (2020): “The Stable Transformation Path,” Working Paper 27731, National Bureau of Economic Research.

Cagliarini, A., and M. Kulish (2013): “Solving Linear Rational Expectations Models with Predictable Structural Changes,” *The Review of Economics and Statistics*, 95(1), 328–336.

Carvalho, V., and X. Gabaix (2013): “The great diversification and its undoing,” *The American Economic Review*, 103(5), 1697–1727.

Comin, D., D. Lashkari, and M. Mestieri (2021): “Structural change with long-run income and price effects,” *Econometrica*, 89(1), 311–374.

Cooley, T. F., and E. C. Prescott (1995): *Frontiers of business cycle research*, vol. 3. Princeton University Press Princeton, NJ.

Da-Rocha, J. M., and D. Restuccia (2006): “The role of agriculture in aggregate business cycles,” *Review of Economic Dynamics*, 9(3), 455–482.

Duarte, M., and D. Restuccia (2010): “The Role of the Structural Transformation in Aggregate Productivity,” *The Quarterly Journal of Economics*, 125(1), 129–173.

——— (2023): “Aggregate Employment and the Rise of Services across Time and Countries,” Discussion paper, University of Toronto.

Duernecker, G., B. Herrendorf, and A. Valentinyi (2021): “The productivity growth slowdown and Kaldor growth facts,” *Journal of Economic Dynamics and Control*, 130, 104200.

Duernecker, G., B. Herrendorf, and Á. Valentinyi (2023): "Structural Change within the Services Sector and the Future of Cost Disease," *Journal of the European Economic Association*, p. jvad030.

Fair, R. C., and J. B. Taylor (1983): "Solution and maximum likelihood estimation of dynamic nonlinear rational expectations models," *Econometrica*, 51(4).

Galesi, A., and O. Rachedi (2019): "Services deepening and the transmission of monetary policy," *Journal of the European Economic Association*, 17(4), 1261–1293.

García-Santana, M., J. Pijoan-Mas, and L. Villacorta (2021): "Investment demand and structural change," *Econometrica*, 89(6), 2751–2785.

Hanoch, G. (1971): "CRESH production functions," *Econometrica: Journal of the Econometric Society*, pp. 695–712.

——— (1975): "Production and demand models with direct or indirect implicit additivity," *Econometrica: Journal of the Econometric Society*, pp. 395–419.

Herrendorf, B., R. Rogerson, and Á. Valentinyi (2013): "Two Perspectives on Preferences and Structural Transformation," *The American Economic Review*, 103(7), 2752–2789.

——— (2014): "Growth and structural transformation," in *Handbook of economic growth*, vol. 2, pp. 855–941. Elsevier.

Herrendorf, B., R. Rogerson, and A. Valentinyi (2021): "Structural change in investment and consumption - a unified analysis," *The Review of Economic Studies*, 88(3), 1311–1346.

Kongsamut, P., S. Rebelo, and D. Xie (2001): "Beyond Balanced Growth," *The Review of Economic Studies*, 68(4), 869–882.

Leon-Ledesma, M., and A. Moro (2020): "The rise of services and balanced growth in theory and data," *American Economic Journal: Macroeconomics*, 12(4), 109–46.

Ljungqvist, L., and T. J. Sargent (2004): "Recursive macroeconomic theory (Second edition)," .

Lucas, R. E. J. (1975): "An equilibrium model of the business cycle," *Journal of political economy*, 83(6), 1113–1144.

Maliar, L., S. Maliar, J. B. Taylor, and I. Tsener (2020): "A tractable framework for analyzing a class of nonstationary Markov models," *Quantitative Economics*, 11(4), 1289–1323.

McConnell, M. M., and G. Perez-Quiros (2000): "Output fluctuations in the United States: What has changed since the early 1980's?," *American Economic Review*, 90(5), 1464–1476.

Moro, A. (2012): "The structural transformation between manufacturing and services and the decline in the US GDP volatility," *Review of Economic Dynamics*, 15(3), 402 – 415.

——— (2015): "Structural Change, Growth, and Volatility," *American Economic Journal: Macroeconomics*, 7(3), 259–94.

Moro, A., and O. Rachedi (2022): "The changing structure of government consumption spending," *International Economic Review*, 63(3), 1293–1323.

Moro, A., and C. Valdes (2021): "Structural Transformation in General Equilibrium," in *New Perspectives on Structural Change*, pp. 156–171. Oxford University Press.

Ngai, L. R., and C. A. Pissarides (2007): "Structural Change in a Multisector Model of Growth," *The American Economic Review*, 97(1), 429–443.

Rogerson, R. (2008): "Structural Transformation and the Deterioration of European Labor Market Outcomes," *Journal of Political Economy*, 116(2), 235–259.

Sato, R. (1975): "The most general class of CES functions," *Econometrica: Journal of the Econometric Society*, pp. 999–1003.

Storesletten, K., B. Zhao, and F. Zilibotti (2019): "Business Cycle during Structural Change: Arthur Lewis Theory from a Neoclassical Perspective," .

Yao, W., and X. Zhu (2021): "Structural change and aggregate employment fluctuations in China," *International Economic Review*, 62(1), 65–100.

Appendix

A The restriction $\sigma = \epsilon_s$ in CLM

This section describes in detail the convenience of setting $\sigma = \epsilon_s$, where σ affects the relative substitution between goods and services and ϵ_s is related to the income elasticity of services. We first note that not imposing this restriction produces estimates that are not statistically different from each other, and that either imposing or not this condition produces a remarkably tight accounting of the long-run data.

To understand the role this assumption plays, consider the Euler equation in CLM:

$$\frac{1}{\beta} = \frac{r_{t+1} + p_{xt+1}(1-\delta)}{p_{xt}} \left(\frac{1 + \Psi_s C_t^{\epsilon_s - 1} p_{st}^{1-\sigma}}{1 + \Psi_s C_{t+1}^{\epsilon_s - 1} p_{st+1}^{1-\sigma}} \right)^{\frac{\sigma}{1-\sigma}} \left(\frac{C_t}{C_{t+1}} \right)^{\tau} \left(\frac{1 + \left(\frac{\sigma - \epsilon_s}{\sigma - 1}\right) \Psi_s C_t^{\epsilon_s - 1} p_{st}^{1-\sigma}}{1 + \left(\frac{\sigma - \epsilon_s}{\sigma - 1}\right) \Psi_s C_{t+1}^{\epsilon_s - 1} p_{st+1}^{1-\sigma}} \right). \tag{6}$$

The Euler equation is used to determine C_{t+1} in our solution algorithm. Given C_t and K_{t+1} (and hence r_{t+1}), equation (6) determinies C_{t+1} via a Bisect method. This method works only when the function studied is strictly monotonic. Next, consider each term containing C_{t+1} in the denominator in the right hand side of equation (6):

$$\underbrace{\left(1 + \Psi_s C_{t+1}^{\epsilon_s - 1} p_{st+1}^{1-\sigma}\right)^{\frac{\sigma}{1-\sigma}}}_{A} \underbrace{C_{t+1}^{\tau}}_{B} \underbrace{\left(1 + \left(\frac{\sigma - \epsilon_s}{\sigma - 1}\right) \Psi_s C_{t+1}^{\epsilon_s - 1} p_{st+1}^{1-\sigma}\right)}_{C}. \tag{7}$$

Term B is clearly increasing in C_{t+1}, and the larger the τ, the larger the derivative. Term A is decreasing since $\epsilon_s > 1$ and $\sigma > 1$. The term C is decreasing as well, since our estimates produce $\sigma < \epsilon_s$ (although not statistically different from each other) when not constraining σ to be equal to ϵ_s. This is a problem for the Bisect method. In particular, one can show that expression (7) is increasing for relatively small values of C_{t+1} and decreasing for larger ones. Thus, the expression has an inverted U-shape as C_{t+1} increases. This not only implies a complication when using the Bisect method, but it also implies that there may be two solutions to equation (6) when solving for C_{t+1}, or none at all.

Setting $\sigma = \epsilon_s$ cancels the term C, making it a constant equal to 1. This is not enough to guarantee that expression (7) is increasing, but it does increase the value of the derivative. A way to guarantee that term B dominates over our range of interest, so that the entire expression is increasing, is to set τ to a relatively large value. In practice, setting $\tau = 3$ suffices to make the expression monotonic. This allows us to solve for C_{t+1} using a Bisect method.

B Estimating equation in CLM

This section derives equation (4). Start with the problem of minimizing expenditures given a level of aggregate consumption:

$$\min_{c_g,c_s}[c_g + p_s c_s] \qquad s.t. \qquad C = \left[\Psi_g^{1/\sigma} C^{\frac{\epsilon_g - 1}{\sigma}} c_g^{\frac{\sigma-1}{\sigma}} + \Psi_s^{1/\sigma} C^{\frac{\epsilon_s - 1}{\sigma}} c_s^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}.$$

The first order conditions to this problem imply

$$c_g = \lambda C^{\frac{1}{\sigma}} \Psi_g^{1/\sigma} C^{\frac{\epsilon_g - 1}{\sigma}} c_g^{\frac{\sigma-1}{\sigma}}, \tag{8}$$

$$p_s c_s = \lambda C^{\frac{1}{\sigma}} \Psi_s^{1/\sigma} C^{\frac{\epsilon_s - 1}{\sigma}} c_s^{\frac{\sigma-1}{\sigma}}, \tag{9}$$

where λ is the Lagrange multiplier. Adding these two equations yields total expenditure in consumption goods:

$$E = c_g + p_s c_s = \lambda C^{\frac{1}{\sigma}} \left[\Psi_g^{1/\sigma} C^{\frac{\epsilon_g - 1}{\sigma}} c_g^{\frac{\sigma-1}{\sigma}} + \Psi_s^{1/\sigma} C^{\frac{\epsilon_s - 1}{\sigma}} c_s^{\frac{\sigma-1}{\sigma}}\right] = \lambda C.$$

Thus, $\lambda = \frac{E}{C}$, so that the Lagrange multiplier is the shadow price of the aggregate consumption good. Inserting this result into equation (8) yields

$$c_g = E^{\sigma} C^{\epsilon_g - \sigma} \Psi_g, \tag{10}$$

and $\eta_g = E^{\sigma-1} C^{\epsilon_g - \sigma} \Psi_g$, where η_g is the share of consumption spent on goods. Simplify equations (8) and (9) to obtain

$$c_g = \lambda^{\sigma} C \Psi_g C^{\epsilon_g - 1}, \tag{11}$$

$$p_s c_s = \lambda^{\sigma} C \Psi_s C^{\epsilon_s - 1} p_s^{1-\sigma}, \tag{12}$$

and divide equation (12) by (11) to obtain:

$$\frac{\eta_s}{\eta_g} = \frac{p_s c_s}{c_g} = \frac{\Psi_s}{\Psi_g} C^{\epsilon_s - \epsilon_g} p_s^{1-\sigma}.$$

Taking logarithms, and noticing from equation (10) that $\log C = \frac{\log \eta_g + (1-\sigma)\log E - \log \Psi_g}{\epsilon_g - \sigma}$ we obtain equation (4).

C Cross-Correlations

In this Appendix we report all cross-correlations among endogenous variables. Tables 6 and 7 display the results.

Table 6: Cross-correlations 1

	C_t,X_t	C_t,h_t	h_t,X_t	C_t,c_{gt}	C_t,c_{st}	X_t,c_{gt}	X_t,c_{st}	c_{gt},c_{st}
Data	0.5637	0.7505	0.8094	0.8694	0.9390	0.3849	0.5502	0.6617
HRV	0.1375	-0.2321	0.8817	0.4428	0.9869	0.8397	0.0228	0.3196
	(0.1463)	(0.1510)	(0.0310)	(0.1989)	(0.0048)	(0.0296)	(0.1540)	(0.1542)
BOP	-0.0351	0.2771	0.0047	0.1788	0.9953	-0.0524	-0.0423	0.1254
	(0.1708)	(0.1627)	(0.1582)	(0.1681)	(0.0013)	(0.1767)	(0.1628)	(0.1705)
CLM	0.0935	-0.3527	0.0107	-0.0040	0.6691	0.8729	-0.5560	-0.7015
	(0.1082)	(0.1555)	(0.1528)	(0.1395)	(0.1053)	(0.0370)	(0.0922)	(0.0608)

Note: The numbers reported in the table are computed as the correlation of the log-deviations of the two variables of interest from their Hodrick-Prescott filtered series. Standard deviations across simulations in parenthesis.

Table 7: Cross-correlations 2

	h_t,c_{gt}	h_t,c_{st}	C_t,L_{st}	C_t,L_{gt}	X_t,L_{gt}	X_t,L_{st}	h_t,L_{gt}	h_t,L_{st}	L_{gt},L_{st}
Data	0.5543	0.7505	0.6914	0.6888	0.8381	0.7273	0.9733	0.9397	0.8497
HRV	0.5276	-0.2321	-0.2647	0.1845	0.4267	0.8501	0.5513	0.9134	0.1807
	(0.0934)	(0.1510)	(0.1460)	(0.1894)	(0.1034)	(0.0512)	(0.1042)	(0.0280)	(0.1443)
BOP	0.8489	0.2771	0.0927	0.5785	0.4530	-0.1355	-0.3070	0.9433	-0.5849
	(0.0533)	(0.1627)	(0.1730)	(0.1356)	(0.2680)	(0.1722)	(0.1633)	(0.0263)	(0.1111)
CLM	0.2259	-0.3527	-0.2619	-0.4083	0.4052	-0.1521	0.8976	0.9793	0.7954
	(0.1615)	(0.1555)	(0.1503)	(0.1048)	(0.1318)	(0.1478)	(0.0326)	(0.0065)	(0.0598)

Note: The numbers reported in the table are computed as the correlation of the log-deviations of the two variables of interest from their Hodrick-Prescott filtered series. Standard deviations across simulations in parenthesis.