NBER WORKING PAPER SERIES

STRUCTURAL CHANGE IN INVESTMENT AND CONSUMPTION: A UNIFIED APPROACH

Berthold Herrendorf
Richard Rogerson
Ákos Valentinyi

Working Paper 24568
http://www.nber.org/papers/w24568

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2018

We thank Julieta Caunedo, V. Kerry Smith, Todd Schoellman and the participants of the ASU Macro Workshop, the London School of Economics, the NBER Growth Group in San Francisco, and the Universities of Köln, Lancaster,Western Ontario, and Windsor for comments and suggestions. Valentinyi thanks the Hungarian National Research, Development and Innovation Office (Project KJS K 124808). All errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w24568.ack

Structural Change in Investment and Consumption: A Unified Approach
Berthold Herrendorf, Richard Rogerson, and Ákos Valentinyi
NBER Working Paper No. 24568
May 2018
JEL No. O11,O14

ABSTRACT

Existing models of structural change typically assume that all of investment is produced in manufacturing. This assumption is strongly counterfactual: in the postwar US, the share of services value added in investment expenditure has been steadily growing and it now exceeds 0.5. We build a new model, which takes a unified approach to structural change in investment and consumption. Our unified approach leads to three new insights: technological change is endogenously investment specific; having constant TFP growth in all sectors is inconsistent with structural change and aggregate balanced growth occurring jointly; the sector with the slowest TFP growth absorbs all resources asymptotically. We also provide empirical support from the postwar US for the first and third insight.

Berthold Herrendorf
W.P. Carey School of Business
Department of Economics
Arizona State University
Tempe, AZ 85287-3806
berthold.Herrendorf@asu.edu

Richard Rogerson
Woodrow Wilson School of
Public and International Affairs
323 Bendheim Hall
Princeton University
Princeton, NJ 08544
and NBER
rdr@princeton.edu

Ákos Valentinyi
Economics, Arthur Lewis Building, 3rd floor
University of Manchester
Oxford Road
Manchester, M13 9PL
United Kingdom
valentinyi.a@gmail.com

1 Introduction

A large, recent literature has proposed extensions of the one-sector growth model in order to jointly study growth and structural change. While the standard framework in the literature allows for multiple consumption goods and structural change among the sectors producing them, it abstracts from structural change within investment and assumes instead that all investment reflects value added from the manufacturing sector.[1] Given that one popular use of these models is to help us understand the decline of the manufacturing sector observed in advanced economies, it seems important to assess the assumption that all investment reflects value added from the manufacturing sector. In this paper we show that the standard assumption is strongly counterfactual and has important implications for the model's predictions. In other words, we show that abstracting from structural change in investment is neither *empirically plausible* nor *theoretically innocuous*.

Our analysis begins by presenting a decomposition of final investment expenditures in the post WWII US into the value added shares coming from the goods sector and the services sector. This analysis is the analogue of the analysis of final consumption expenditure in Herrendorf et al. (2013). Two key findings emerge. First, although it is true that the share of goods value added has always been much higher in investment expenditure than in consumption expenditure, we show that the standard assumption is squarely at odds with the data. Specifically, the share of services value-added in investment expenditure is large, and in fact now exceeds the share of goods value added in investment expenditure. Second, and more importantly, there is structural change within the investment sector and it features the same qualitative patterns as structural change within the consumption sector: the expenditure share of services has increased at the same time that the relative price of services has increased.

Motivated by these facts, we develop a general equilibrium model of growth and structural change. Our framework can accommodate various levels of disaggregation, but, to best highlight the key economic interactions, our core analysis focuses on the simplest two-by-two-by-two structure, featuring two final expenditure categories (investment and consumption), two underlying sectors that produce value added (goods and services), and two factor inputs in the production of value added (capital and labor). Final investment and final consumption are produced by combining value added from the goods and services sectors which are in turn produced by labor and capital. Notably, we take a unified approach and treat the production of consumption and investment in a symmetric manner by assuming that each is a CES aggregator of goods and services value added, with possibly different weights and elasticities of substitution.

Having developed our new model, we proceed to examine the properties of its equilibrium.

[1]This practice started with the early models of growth and structural change by Echevarria (1997) and Kongsamut et al. (2001). See Herrendorf et al. (2014) for a review of the literature on structural change that has emerged since then. An important recent exception is Garcia-Santana et al. (2016), which we discuss in more detail below.

We show that our framework allows for a two-step procedure for analyzing structural change and balanced growth. Specifically, one can first analyze balanced growth focusing purely on aggregate consumption and investment, and then exam structural change after balanced growth has been established. Key to this result is establishing the importance of a concept we call *effective TFP* for the investment sector. This result generalizes the result in Herrendorf et al. (2014) to a setting that also allows for structural change in investment. We provide necessary and sufficient conditions for a generalized balanced growth path to exist.

Structural change within each of consumption and investment is dictated by a standard force: changes in relative prices combined with a less than unitary elasticity implies that expenditure shares increase for the input whose relative price increases. We call this the intensive margin of structural change. If investment and consumption have different value added shares of goods and services, aggregate structural change can occur also via an extensive margin. The extensive margin results from changes in the mix of consumption and investment, which lead to structural change because investment and consumption have different value added mixes of goods and services value added. We show that along a balanced growth path there is no adjustment along the extensive margin, so that all structural change comes from the intensive margin.

Taking a unified approach to structural change in investment and consumption leads to distinct conclusions along several key dimensions. In particular, three key insights emerge. First, for empirically plausible parameter values, technological change is endogenously investment specific. Second, we derive a necessary and sufficient condition for the existence of a balanced growth path and show that constant (but possibly different) growth in each of the three TFP terms is inconsistent with structural change and aggregate balanced growth occurring jointly. Third, the sector with the slowest TFP growth absorbs all resources asymptotically.

The CES aggregators in the consumption and investment sectors of our simple framework are admittedly somewhat specialized, and so there may be concern about how well our specification performs quantitatively. We show that it can quantitatively capture most of the salient features of structural change in the US data in the post WWII period, with the best fit coming from a specification in which there is little scope for substitution between goods and services in the production of both investment and consumption. This result extends the earlier analysis of Herrendorf et al. (2013) to the case of structural change within the investment sector. This specification captures structural change within investment remarkably well.

We also carry out an empirical assessment of our theoretical condition that is necessary for balanced growth. Our theoretical condition places a non-linear restriction on the evolution of the sectoral TFPs in our model. We estimate the sectoral TFP growth rates using standard growth accounting methods and evaluate the extent to which our theoretical condition holds and the role of changes in each of the sectoral TFPs. We find that our theoretical condition holds approximately in the data. Interestingly, however, it holds despite the fact that the growth rates of the sectoral TFP terms vary quite dramatically over time. This suggests that, contrary to common practice, constant growth of sectoral TFPs is not a natural restriction to impose on the parameters in the context of balanced growth in multi-sector models. We also find that quantitatively most of the investment-specific technological change arises endogenously. We conclude from this finding that our first insight ("technological change is endogenously investment biased") is far more than a theoretical curiosity.

Our work is closely related to a recent paper by Garcia-Santana et al. (2016). They also begin with the observation that investment is produced by a combination of goods and services and that goods represent a higher share of value added for investment than for consumption. Despite this similarity, the papers focus on distinct issues. We focus on analytically establishing the existence and properties of a generalized balanced growth path with structural change in a general equilibrium setting. Structural change in our analysis arises entirely from the *intensive* margin, because we show that along a generalized balanced growth path the extensive margin is not operational. In contrast, they focus on transition dynamics in a partial equilibrium setting, and establish empirically that the *extensive* margin is a key driver of structural change away from the balanced growth path. Specifically, they study how movements in the investment-to-GDP ratio, for example during an investment boom, produce hump-shaped movements in the value added share of the manufacturing sector in a large sample of countries.[2]

An outline of the paper follows. In the next section we present the key facts from the US time series data. Section 3 presents our model and Section 4 characterizes the key properties of the competitive equilibrium. Section 5 studies the features of structural change along a GBGP. Section 6 derives the three insights that result from our unified approach. Section 7 examines key aspects of the model and its theoretical properties from an empirical perspective. Section 8 discusses extending the analysis to consider three different categories of investment expenditure, and Section 9 concludes.

2 Evidence

In this section we offer a unified analysis of the structural-change facts for both final investment and final consumption for the US over the post WWII period. This serves to both complement existing presentations of the stylized facts of structural change and motivate the framework that we develop in the next section.[3] The basic strategy is to combine US industry data from WORLD KLEMS with the annual input-output tables from the BEA and to then decompose

[2]We will point further, more subtle differences between our work and their work as we go along.

[3]Our evidence complements that presented in Garcia-Santana et al. (2016) for 40 developed countries covering a recent time period. Whereas we plot the sectoral shares for a single country (the US) over a long period, they pool short time series data for many countries to characterize how sectoral shares vary with GDP per capita.

Figure 1: Sector shares in consumption and investment

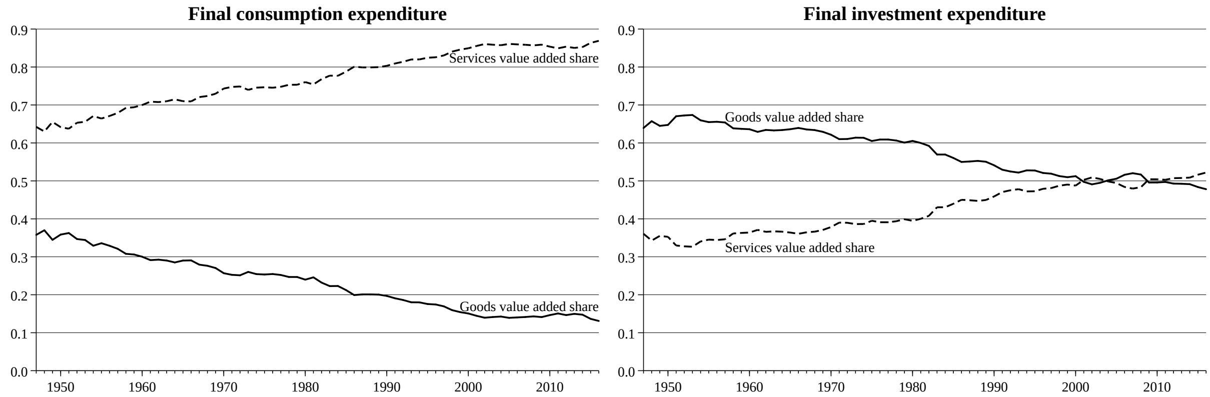

final expenditures into its value added components. One can implement this decomposition at various levels of disaggregation, but to best highlight the novel implications of our analysis we consider two final expenditure categories – consumption and investment – and two value-added components – goods and services. In a later section we present evidence for an alternative decomposition in which we consider three categories of investment – structures, equipment and intellectual property product.

In generating these decompositions we define the goods sector to consist of agriculture, construction, manufacturing, mining, and public utilities.[4] Services consists of the remaining industries – business services, government, personal services, transportation, wholesale and retail trade. Our analysis in the following sections focuses on the case of a closed economy. To connect the closed economy model to the data requires allocating net exports between consumption and investment. Because net exports are not that large, the rule for allocating them is not of first-order significance. In what follows we allocate all of net exports to consumption. One benefit of this choice is that the notion of investment and capital in the model will correspond to the notion of investment and capital as measured in the data.

We decompose each final expenditure category into the value added from different sectors following the methodology developed in Herrendorf et al. (2013), which involves the use of input–output relationships and total requirement matrices. Note that while in Herrendorf et al. (2013) we decomposed final consumption expenditure into the value added from agriculture, manufacturing, and services, here we decompose final investment expenditure and final consumption expenditure into the value added components produced in the goods and services sectors.

[4]Much of the structural change literature considers three sectors: agriculture, non-agricultural goods (typically referred to as manufacturing) and services. Because we focus on the US during the post World War II period when agriculture is relatively unimportant, we have chosen to combine agriculture and manufacturing into a single goods producing-sector. We hope that this facilitates exposition.

Figure 2: Relative prices

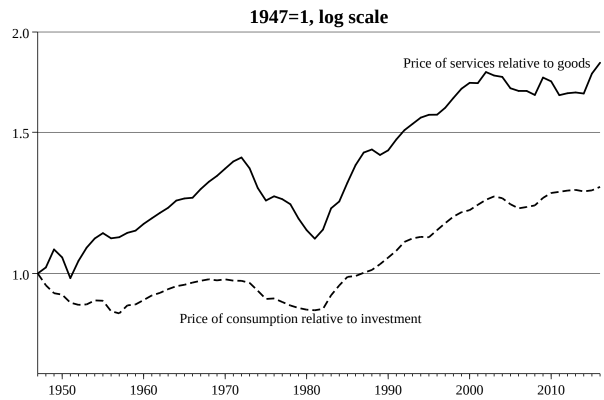

Figure 1 shows the key facts for this two–by–two decomposition. While investment has a significantly higher goods-valued-added share than does consumption, we see that both consumption and investment exhibit an increase in the services-value-added share and a decrease in the goods-value-added share.

We stress three key properties relative to the existing literature on structural change. First, assuming that investment is produced entirely by the goods sector is strongly at odds with the data; in fact, by the end of the sample the goods value added share is less than the services value added share in investment.[5] Second, the value added shares exhibit important changes over time in both investment and consumption. This suggests that any analysis of structural change at the aggregate level needs to confront the reality that structural change occurs both within the investment–producing sector and the consumption–producing sector. Third, the value added shares of goods differ significantly between investment and consumption, suggesting that it is important that consumption and investment be modelled separately.

Because relative prices will play an important role in the analysis to come, we also present evidence on two key relative price movements over time. In particular, Figure 2 displays time series evidence on the prices of services relative to goods and consumption relative to investment. The Figure reveals two key secular trends that are familiar from the existing literature. First, there has been a marked increase in the price of services relative to goods. The somewhat unusual behavior of this relative price in the 1970s is driven by a dramatic spike in the prices of agricultural products and oil during the early 1970s. While this suggests that a more detailed analysis might warrant further disaggregation, the somewhat anomalous behavior of the 1970s should not distract us from the clear secular trend over the entire postwar period. Second, there is also a marked increase in the price of consumption relative to investment. Notably, the

[5]As noted in Herrendorf et al. (2013), a different but related problem with the standard assumption is that in recent years, the value of US investment has exceeded the value added produced in the goods sector.

behavior is quite distinct between the pre- and post-1980 periods, with little trend in the first subperiod and a marked positive trend in the second subperiod.

3 Model

We build a multi-sector extension of the standard one-sector neoclassical growth model, formulated in continuous time. Motivated by the presentation in the previous section, our approach is to start with sectoral valued-added production functions and to model the production for final expenditure categories by aggregating the sectoral value added.

Specifically, we assume that goods and services value added (denoted by Y_{gt} and Y_{st}, respectively) are each produced according to Cobb-Douglas production functions with the same capital shares but potentially different rates of technological progress:

$$Y_{jt} = A_{jt} K_{jt}^{\theta} L_{jt}^{1-\theta}, \quad j \in \{g, s\},$$

where $\theta \in (0, 1)$ and the A_{jt} represent exogenous technological progress. The assumption that the underlying production functions are Cobb-Douglas with identical capital shares is common in the literature on structural change, as it allows for the existence of a balanced growth path.[6] Moreover, Herrendorf et al. (2015) show that this case also does a reasonable job of replicating the labor allocation across sectors in the postwar US.

The outputs of the goods and services sectors are in turn combined to produce final consumption and final investment using CES aggregators:

$$C_t = \left(\omega_c^{\frac{1}{\varepsilon_c}} C_{gt}^{\frac{\varepsilon_c - 1}{\varepsilon_c}} + (1 - \omega_c)^{\frac{1}{\varepsilon_c}} C_{st}^{\frac{\varepsilon_c - 1}{\varepsilon_c}} \right)^{\frac{\varepsilon_c}{\varepsilon_c - 1}}, \tag{1}$$

$$X_t = A_{xt} \left(\omega_x^{\frac{1}{\varepsilon_x}} X_{gt}^{\frac{\varepsilon_x - 1}{\varepsilon_x}} + (1 - \omega_x)^{\frac{1}{\varepsilon_x}} X_{st}^{\frac{\varepsilon_x - 1}{\varepsilon_x}} \right)^{\frac{\varepsilon_x}{\varepsilon_x - 1}}, \tag{2}$$

where $\varepsilon_j \in [0, \infty)$ is the elasticity of substitution between goods and services in the production of final expenditure category j and $\omega_j \in [0, 1]$ determines the relative weight of inputs from the goods sector into the production of final expenditure category j, for $j \in \{c, x\}$. The standard case in which investment is entirely produced in the goods sector is captured as the special case in which $\omega_x = 1$.[7] A_{xt} represents exogenous investment–specific technological change. Any common TFP component for these two aggregators is equivalent to higher TFP in the goods and services production functions and so is not separately identified. For this reason we

[6] In a different context, Acemoglu and Guerrieri (2008) and Alvarez-Cuadrado et al. (2017) study a model in which the two technologies have different capital shares and show that a balanced growth path exists only asymptotically.

[7] Note that ω_j and $1 - \omega_j$ are raised by $1/\varepsilon_j$ so as to ensure that the limit for $\varepsilon_j \to 0$ is a Leontief with weights ω_j and $1 - \omega_j$.

normalize the TFP level for the consumption aggregator to unity. We assume that the growth rates of the three TFP terms are all bounded.

While these CES specifications are analytically convenient, we show later that they also do a reasonable job of capturing the empirical patterns presented in the previous section. A key feature of our model is that we treat the production of consumption and investment symmetrically.

We make two remarks concerning the specification for C_t. First, our consumption aggregator is homothetic. The literature has argued that non-homotheticities account for at least part of the structural change in consumption and, as we will see in Section 7, our specification will not be able to account for all of the structural change within consumption.[8] We nonetheless adopt a homothetic aggregator for consumption because it allows us to better focus on the novel implications of our formulation of investment. Importantly, our specification does a very good job of capturing structural change within final investment expenditure.

Second, it does not matter whether we think of the household buying goods and services in the market and combining them itself to produce C_t, or alternatively, that a stand-in firm purchases goods and services, combines them into the aggregate C_t, and then sells the aggregate consumption good to the household. The advantage of the latter formulation is that it provides an explicit market price for the consumption aggregate, which will be useful for the analysis that follows. We will therefore view relationship (1) as a production function for a firm that operates in the market.

There is an infinitely lived representative household with preferences represented by the utility function

$$\int_0^\infty e^{-\rho t} \log(C_t) dt,$$

where $\rho > 0$ is the discount rate. The household is endowed with one unit of time at each instant, which is supplied inelastically, and a positive initial capital stock, $K_0 > 0$.

Capital depreciates at rate $\delta \in (0, 1]$, so the law of motion for capital is given by:

$$\dot{K}_t = X_t - \delta K_t.$$

Capital and labor are freely mobile across sectors. Feasibility then requires:

$$K_{gt} + K_{st} \leq K_t, \qquad L_{gt} + L_{st} \leq 1,$$
$$C_{gt} + X_{gt} \leq Y_{gt}, \qquad C_{st} + X_{st} \leq Y_{st}.$$

[8] Several papers have recently studied the role that income effects play for structural change. In addition to the early paper by Kongsamut et al. (2001), recent examples include Herrendorf et al. (2013), Boppart (2014), Comin et al. (2015), and Duernecker et al. (2017b).

4 Equilibrium

We study the competitive equilibrium for the above economy. We assume four representative firms, one for each of goods, services, consumption and investment, and assume that the household accumulates capital and rents it to the firms. At each point in time there will be six markets: four markets for the firm outputs and two markets for production factors. Rental prices for capital and labor are denoted by R_t and W_t respectively, the prices for goods and services are denoted by P_{gt} and P_{st} , respectively, and the prices of final consumption and final investment are denoted by P_{ct} and P_{xt}. We normalize the price of the final investment good to one in each period. Arbitrage implies that the interest rate will equal $R_t - \delta$. Given that the equilibrium concept is standard, we do not provide a formal definition of equilibrium.[9]

In the remainder of this section, we provide a partial characterization of the equilibrium. We first derive analytic expressions for the prices of the four outputs in terms of model primitives. As a by product of the first step we also generate expressions for relative expenditure shares on inputs in each of the final sectors in terms of model primitives. Lastly, we derive alternative representations to characterize production in equilibrium that provide a connection between our model and more standard two-sector versions of the growth model.

4.1 Output Prices

We start with the first-order conditions for capital and labor for the two firms producing goods and services. For $j \in \{g, s\}$ these are given by:

$$\theta P_{jt} A_{jt} K_{jt}^{\theta-1} L_{jt}^{1-\theta} = R_t, \tag{3}$$

$$(1-\theta) P_{jt} A_{jt} K_{jt}^{\theta} L_{jt}^{-\theta} = W_t. \tag{4}$$

Taking the ratio of the two first-order conditions for a given sector j and rearranging gives:

$$\frac{K_{jt}}{L_{jt}} = \frac{\theta}{1-\theta} \frac{W_t}{R_t}. \tag{5}$$

It follows that capital-labor ratios are equalized. Given that aggregate labor is one, it follows that $K_{jt}/L_{jt} = K_t$ for all t. Using this fact, the two first-order conditions for capital imply that relative prices are the inverse of relative TFPs:

$$\frac{P_{gt}}{P_{st}} = \frac{A_{st}}{A_{gt}} \tag{6}$$

[9]Because the equilibrium allocation of our model is efficient, we could instead study the planner problem. But since the evidence on relative prices in Figure 2 provides important information regarding parameter values, it is useful to study the competitive equilibrium directly.

This is a standard result in the structural change literature when the sector production functions are Cobb Douglas with the same capital-share parameter.

Next we derive expressions for P_{ct} and P_{xt} in terms of P_{gt} and P_{st}. To do this we use the fact that constant-returns-to-scale production functions imply that profits must equal zero in a competitive equilibrium. It follows that P_{ct} must equal the minimum cost to produce a unit of consumption, and similarly that P_{xt} (which is normalized to one) must equal the minimum cost of producing a unit of investment. Straightforward calculation yields:

$$P_{ct} = \left(\omega_c P_{gt}^{1-\varepsilon_c} + (1-\omega_c) P_{st}^{1-\varepsilon_c}\right)^{\frac{1}{1-\varepsilon_c}}, \tag{7}$$

$$1 = \frac{\left(\omega_x P_{gt}^{1-\varepsilon_x} + (1-\omega_x) P_{st}^{1-\varepsilon_x}\right)^{\frac{1}{1-\varepsilon_x}}}{A_{xt}}. \tag{8}$$

The three equations (6), (7), and (8) allow us to fully characterize the three prices P_{gt}, P_{st} and P_{ct} in terms of primitives. Specifically, equations (6) and (8) are two equations in the two unknowns P_{gt} and P_{st}.[10] Straightforward algebra yields:

$$P_{gt} = \frac{A_{xt}\left(\omega_x A_{gt}^{\varepsilon_x-1} + (1-\omega_x) A_{st}^{\varepsilon_x-1}\right)^{\frac{1}{\varepsilon_x-1}}}{A_{gt}}, \tag{9}$$

$$P_{st} = \frac{A_{xt}\left(\omega_x A_{gt}^{\varepsilon_x-1} + (1-\omega_x) A_{st}^{\varepsilon_x-1}\right)^{\frac{1}{\varepsilon_x-1}}}{A_{st}}. \tag{10}$$

Substituting these equations into equation (7) gives:

$$P_{ct} = \frac{A_{xt}\left(\omega_x A_{gt}^{\varepsilon_x-1} + (1-\omega_x) A_{st}^{\varepsilon_x-1}\right)^{\frac{1}{\varepsilon_x-1}}}{\left(\omega_c A_{gt}^{\varepsilon_c-1} + (1-\omega_c) A_{st}^{\varepsilon_c-1}\right)^{\frac{1}{\varepsilon_c-1}}}. \tag{11}$$

4.2 Expenditure Shares

The cost minimization problems for the production of final consumption and final investment also generate standard expressions for relative expenditures on inputs in each of the two final expenditure sectors:

$$\frac{P_{gt} C_{gt}}{P_{st} C_{st}} = \frac{\omega_c}{1-\omega_c}\left(\frac{P_{gt}}{P_{st}}\right)^{1-\varepsilon_c}, \tag{12}$$

$$\frac{P_{gt} X_{gt}}{P_{st} X_{st}} = \frac{\omega_x}{1-\omega_x}\left(\frac{P_{gt}}{P_{st}}\right)^{1-\varepsilon_x}. \tag{13}$$

[10] Note that the zero-profit condition for the investment sector can be rewritten as an expression linking the price of goods and the relative price of goods to services. Since this relative price is determined by relative TFPs, this equation gives us the price of goods in terms of primitives.

These expressions describe the nature of structural change within consumption and investment as a function of the relative price of goods to services and the elasticities of substitution. Given data on relative prices and relative expenditure shares, these expressions can be used to infer values for the ω_j and the ε_j. We carry out this exercise in Section 7.

Combined with our previous result about the price of goods relative to services in equation (2), we can also express expenditure shares purely as a function of model primitives:

$$\begin{aligned}\frac{P_{gt}C_{gt}}{P_{st}C_{st}} &= \frac{\omega_c}{1-\omega_c}\left(\frac{A_{st}}{A_{gt}}\right)^{1-\varepsilon_c} \\ \frac{P_{gt}X_{gt}}{P_{st}X_{st}} &= \frac{\omega_x}{1-\omega_x}\left(\frac{A_{st}}{A_{gt}}\right)^{1-\varepsilon_x}\end{aligned}$$

4.3 Alternative Representations of Production

Our formulation of technology involved a two-level structure in which we started with value-added production functions for goods and services, and then produced final-expenditure categories by combining valued added from the goods and services sectors. A common alternative formulation of the two-sector growth model is to directly posit value added productions for each of consumption and investment. More concretely, this approach would assume:

$$\begin{aligned}C_t &= \tilde{A}_{ct}K_{ct}^{\theta}L_{ct}^{1-\theta} \\ X_t &= \tilde{A}_{xt}K_{xt}^{\theta}L_{xt}^{1-\theta}\end{aligned}$$

A simple prediction of this model is that in equilibrium, the price of consumption relative to investment is equal to $\tilde{A}_{xt}/\tilde{A}_{ct}$, so that relative TFP growth of the two sectors can be inferred directly from data on relative price changes. Additionally, a balanced growth path exists as long as $\tilde{A}_{xt}$ grows at a constant rate.

This formulation abstracts from the process of structural transformation, a feature that is the focus of our analysis. However, because the analysis of balanced growth is well established in this setting it is instructive to provide a connection between this formulation and our two-level formulation, a task that we turn to in this subsection.

To develop this connection we derive the mappings between equilibrium inputs of capital and labor to equilibrium outputs of consumption and investment. Intuitively, given equilibrium prices, we can solve for the optimal input mix of goods and services for each of C and X. Because we know the equilibrium capital-labor ratios in each of goods and services, we can then infer the relative inputs of capital and labor used to produce the optimal mix of goods and services in each of C and X. The next proposition characterizes this mapping.

Proposition 1 *Along any equilibrium path, the following hold:*

$$X_t = \mathcal{A}_{xt} K_{xt}^{\theta} L_{xt}^{1-\theta}, \tag{14}$$

$$C_t = \mathcal{A}_{ct} K_{ct}^{\theta} L_{ct}^{1-\theta}, \tag{15}$$

where

$$\mathcal{A}_{xt} \equiv A_{xt}\left(\omega_x A_{gt}^{\varepsilon_x-1} + (1-\omega_x) A_{st}^{\varepsilon_x-1}\right)^{\frac{1}{\varepsilon_x-1}}, \quad L_{xt} \equiv \frac{X_{gt}}{A_{gt}K_t^{\theta}} + \frac{X_{st}}{A_{st}K_t^{\theta}}, \qquad K_{xt} \equiv K_t L_{xt}, \tag{16}$$

$$\mathcal{A}_{ct} \equiv \left(\omega_c A_{gt}^{\varepsilon_c-1} + (1-\omega_c) A_{st}^{\varepsilon_c-1}\right)^{\frac{1}{\varepsilon_c-1}}, \qquad L_{ct} \equiv \frac{C_{gt}}{A_{gt}K_t^{\theta}} + \frac{C_{st}}{A_{st}K_t^{\theta}}, \qquad K_{ct} \equiv K_t L_{ct}. \tag{17}$$

Proof. See the Appendix.

The two expressions (14) and (15) have the appearance of production functions, with the $\mathcal{A}_{jt}$ as TFPs. But because they are relations that hold only in equilibrium we will refer to them as *pseudo production functions* and we will refer to the $\mathcal{A}_{jt}$ as *effective TFPs*. Our model posits Cobb-Douglas production functions at the sector level, whose outputs are aggregated in non-linear ways to produce final consumption and final investment. Interestingly, these pseudo-production functions show that the Cobb-Douglas property is preserved with regard to inputs of labor and capital and that the non-linear aggregation only affects the expressions for effective TFPs.

Note that given the definition of the effective TFPs, equation (11) is equivalent to:

$$P_{ct} = \frac{\mathcal{A}_{xt}}{\mathcal{A}_{ct}},$$

which is the analogue of the result that holds in the more standard formulation of the two-sector growth model. Importantly, and a point that we will return to later, the ratio of effective TFPs is a non-linear function of all three TFPs in our model, and is not the same as A_{xt}, the TFP for production of investment relative to consumption.

It is also instructive to derive another pseudo production function, this one for aggregate final output, where we define aggregate final output as measured in units of final investment: $Y_t = X_t + P_{ct}C_t$.[11] In the context of the standard two sector model it is easy to show that one obtains $Y_t = \tilde{A}_{xt}K_t^{\theta}$. See, for example the derivations in Herrendorf et al. (2014). We next develop the analogous relationship for our model. Because payments to inputs will exhaust income for each of the investment- and consumption-producing firms, we can also write:

$$Y_t = P_{gt}X_{gt} + P_{st}X_{st} + P_{gt}C_{gt} + P_{st}C_{st} = P_{gt}(X_{gt} + C_{gt}) + P_{st}(C_{st} + X_{st}).$$

[11] Duernecker et al. (2017a) stress that this is not the way in which output is defined in NIPA and that one needs to be careful when mapping model output into data output. We nonetheless use this definition of output because it allows us to analytically characterize the equilibrium path of the model as a generalized balanced growth path.

Feasibility requires that $X_{gt} + C_{gt} = Y_{gt}$ and $X_{st} + C_{st} = Y_{st}$. Substituting these conditions and using the fact that $P_{gt}A_{gt} = P_{st}A_{st}$ gives:

$$\begin{aligned} Y_t &= P_{gt}A_{gt}K_{gt}^{\theta}L_{gt}^{1-\theta} + P_{st}A_{st}K_{st}^{\theta}L_{st}^{1-\theta} \\ &= P_{gt}A_{gt}\left(K_{gt}^{\theta}L_{gt}^{1-\theta} + K_{st}^{\theta}L_{st}^{1-\theta}\right). \end{aligned}$$

Using $K_{gt}/L_{gt} = K_{st}/L_{st} = K_t$ and $L_{gt} + L_{st} = 1$, we have:[12]

$$\begin{aligned} Y_t &= P_{gt}A_{gt}\left[\left(\frac{K_{gt}}{L_{gt}}\right)^{\theta} L_{gt} + \left(\frac{K_{st}}{L_{st}}\right)^{\theta} L_{st}\right] \\ &= P_{gt}A_{gt}K_t^{\theta}\left(L_{gt} + L_{st}\right) \\ &= P_{gt}A_{gt}K_t^{\theta}. \end{aligned}$$

Using the previously derived expression for P_{gt} in terms of primitives from equation (9) gives:

$$Y_t = A_{xt}\left(\omega_x A_{gt}^{\varepsilon_x-1} + (1-\omega_x)A_{st}^{\varepsilon_x-1}\right)^{\frac{1}{\varepsilon_x-1}} K_t^{\theta} = \mathcal{A}_{xt}K_t^{\theta}. \tag{18}$$

where $\mathcal{A}_{xt}$ is the effective TFP in the investment sector as defined earlier. Note that the aggregate input of labor does not feature explicitly in equation (18) because it equals one.

This expression is the natural generalization of the relation that one obtains in benchmark two-sector models with exogenous sector-specific technical change formulated directly with value added production functions for consumption and investment. As noted above, the key difference is that in our framework the level of effective TFP in the investment sector is a non-linear function of all three fundamental TFP levels in our economy and so is not simply equal to $\tilde{A}_{xt}$.

Recalling that we have normalized the price of investment to unity, it is straightforward to show that in competitive equilibrium, factor prices are given by:

$$\begin{aligned} R_t &= \theta\mathcal{A}_{xt}K_t^{\theta-1} \\ W_t &= (1-\theta)\mathcal{A}_{xt}K_t^{\theta} \end{aligned}$$

5 Structural Change and Balanced Growth

We are now ready to explore the implications of our model for structural change and aggregate balanced growth.

[12] Note that since $P_{gt}A_{gt} = P_{st}A_{st}$ the last step in this derivation could also be written using the price of services and TFP in the service sector.

5.1 Measuring Structural Change

The dimension of structural change that concerns us is the composition of aggregate economic activity across goods and services. This can be measured either as changes in sectoral value added shares or changes in sectoral employment shares. In our model these two measures are identical in equilibrium, a property shared by many models in the literature. To see this, note that the nominal value added shares for goods and services, which we denote as V_{jt} ($j \in \{g, s\}$), are given by:

$$V_{jt} \equiv \frac{P_{jt}Y_{jt}}{Y_t} = \frac{P_{jt}Y_{jt}}{P_{gt}Y_{gt} + P_{st}Y_{st}}.$$

The second equality derives from the fact that with constant returns to scale, payments to factors of production exhaust all income. Substituting the sectoral production functions and using the facts that in equilibrium $P_{gt}A_{gt} = P_{st}A_{st}$ and $K_{gt}/L_{gt} = K_{st}/L_{st} = K_t$, we have:

$$V_{jt} = \frac{P_{jt}A_{jt}K_t^{\theta}L_{jt}}{P_{gt}A_{gt}K_t^{\theta}L_{gt} + P_{st}A_{st}K_t^{\theta}L_{st}} = \frac{L_{jt}}{L_{gt} + L_{st}} = L_{jt}.$$

In other words, sectoral nominal value added shares equal sectoral employment shares and we can restrict our attention to characterizing the former.

5.2 Intensive and Extensive Margins of Structural Change

Given we have only two sectors, it is sufficient to focus on the behavior of V_{st}. Using $Y_{st} = C_{st} + X_{st}$ and carrying out some simple manipulations gives:

$$\begin{aligned} V_{st} &= \frac{P_{ct}C_t}{Y_t}\frac{P_{st}C_{st}}{P_{ct}C_t} + \frac{X_t}{Y_t}\frac{P_{st}X_{st}}{X_t} \\ &= \frac{P_{ct}C_t}{Y_t}\frac{P_{st}C_{st}}{P_{gt}C_{gt} + P_{st}C_{st}} + \frac{X_t}{Y_t}\frac{P_{st}X_{st}}{P_{gt}X_{gt} + P_{st}X_{st}} \end{aligned}$$

This can be written as:

$$V_{st} = \underbrace{\frac{P_{ct}C_t}{Y_t}}_{\textbf{extensive margin}} \underbrace{\frac{1}{(P_{gt}C_{gt})/(P_{st}C_{st}) + 1}}_{\textbf{intensive margin}} + \underbrace{\frac{X_t}{Y_t}}_{\textbf{extensive margin}} \underbrace{\frac{1}{(P_{gt}X_{gt})/(P_{st}X_{st}) + 1}}_{\textbf{intensive margin}}.$$

We previously solved for the ratio of the expenditure shares of goods to services in terms of primitives:

$$\frac{P_{gt}C_{gt}}{P_{st}C_{st}} = \frac{\omega_c}{1-\omega_c}\left(\frac{A_{st}}{A_{gt}}\right)^{1-\varepsilon_c}, \tag{19}$$

$$\frac{P_{gt}X_{gt}}{P_{st}X_{st}} = \frac{\omega_x}{1-\omega_x}\left(\frac{A_{st}}{A_{gt}}\right)^{1-\varepsilon_x}. \tag{20}$$

Hence, within each final expenditure sector, structural change is driven by the standard forces that operate in the presence of uneven technological change: if $\widehat{A_{gt}} > \widehat{A_{st}}$ and $\varepsilon_c, \varepsilon_c \in [0, 1)$, then the expenditure share on goods is decreasing and the expenditure share on services in both final consumption and final investment. We call this the intensive margin of structural change, as it operates within each final expenditure category. Note that although the intensive margin of structural change has the same driving force within each final expenditure category, as long as ω_x and ω_c are different, the magnitude of intensive margin structural change will differ across the two categories.

From the perspective of the overall economy, structural change is also potentially influenced by the extensive margin that arises from changes in the composition of final expenditure between investment and consumption. If the final expenditure shares are constant over time, then the above effects imply that labor is moving systematically from the goods-producing sector to the services-producing sector. But if there is a change in final expenditure shares between consumption and investment, then one cannot infer the nature of structural change without additional assumptions. This is because expenditure could be reallocated to investment – the sector with a higher goods share – at a sufficiently high rate so that the overall goods share could rise even though it declines within each sector.

In the next subsection we show that along a generalized balanced growth path (GBGP), the final expenditure shares are indeed constant so that only the intensive margin operates and the nature of structural change along such a path can be inferred without additional assumptions. In contrast, Garcia-Santana et al. (2016) focus on changes in the extensive margin associated with dynamics away from the balanced growth path. Specifically, they show that investment booms associated with high growth episodes can give rise to hump-shaped manufacturing shares.

5.3 Generalized Balanced Growth

We previously derived properties that will hold along any equilibrium path. In this subsection we examine the properties of our model along a GBGP. As mentioned in the introduction, balanced growth is too strict an equilibrium concept to impose in settings with structural change, since sectoral labor inputs cannot feasibly grow at a constant yet different rates forever. The term *generalized* balanced growth relaxes balanced growth by requiring that *aggregate* variables grow at constant rates or remain constant without restricting the growth of *sectoral* variables.

We adopt the definition of Kongsamut et al. (2001), who defined a GBGP as an equilibrium path along which the rental price of capital is constant.[13]

To this point, our analysis has only made use of equilibrium conditions from the production side of the economy. To characterize the properties of a GBGP, we will also need to use the equilibrium conditions that arise from the household side of the equilibrium. The household seeks to maximize the present discounted integral of utility subject to its budget equation:

$$E_t + \dot{K}_t + \delta K_t = R_t K_t + W_t, \tag{21}$$

where $E_t \equiv P_{ct} C_t$ denotes consumption expenditure. Using "hats" to denote growth rates, the Euler equation and the transversality condition for the household can be written as:

$$\widehat{E_t} = R_t - \delta - \rho, \tag{22}$$

$$\lim_{t\to\infty} e^{-\rho t} \frac{K_t}{E_t} = 0. \tag{23}$$

The first proposition establishes that if a GBGP exists, it must be the case that K_t, X_t, Y_t, and E_t all grow at the same constant rate, which is determined by the right-hand side of the Euler equation. Note that whereas both K_t and X_t are physical quantities, the other two terms – Y_t and E_t – include both prices and quantities.

Proposition 2 *If a GBGP exists and R is the rental rate along the GBGP, then it must be that*

$$\widehat{K_t} = \widehat{X_t} = \widehat{Y_t} = \widehat{E_t} = R - \delta - \rho.$$

Proof. See the Appendix.

Having established the properties of a GBGP in Proposition 2, we can now revisit the issue of structural change along the GBGP. As noted in the previous subsection, although knowing the direction of relative price changes for goods and services together with the elasticity of substitution between goods and services in production is sufficient to determine the intensive margins of structural change within consumption and investment individually, changes in the composition of final output between consumption and investment could create an opposing effect via the extensive margin. But what we have just shown is that along a GBGP the composition of final expenditure is constant, and thus there is no opposing effect from composition changes. It follows:

Proposition 3 *If a GBGP exists, $\varepsilon_x, \varepsilon_c \in (0, 1)$, and $\widehat{A_{gt}} > \widehat{A_{st}}$, then the shares of the services sector in total value added and total employment will increase over time. The value added share of the services sector will also increase for both final consumption and final investment.*

[13] This is equivalent to a constant real interest rate, which equals the rental price of capital minus the depreciation rate.

The previous propositions characterized what happens along a GBGP assuming that such a path exists. The next proposition provides necessary and sufficient conditions for the existence of a GBGP.

Proposition 4 *A GBGP exists if and only if $\mathcal{A}_{xt}$ grows at a constant rate.*

To show the claims of this proposition, we use the pseudo aggregate production function derived previously:

$$Y_t = \mathcal{A}_{xt} K_t^{\theta}.$$

Taking the derivative with respect to capital, the requirement that the rental price of capital be constant becomes:

$$R_t = R = \theta \mathcal{A}_{xt} K_t^{\theta-1}. \tag{24}$$

We previously showed that along a GBGP, K_t necessarily grows at a constant rate. It follows that if R_t is constant then $\mathcal{A}_{xt}$ must also grow at a constant rate. Conversely, suppose that $\mathcal{A}_{xt}$ grows at a constant rate. Given that K_t must grow at rate $R - \delta - \rho$ there is a unique value of R that is consistent with a constant rental rate given the constant growth rate of $\mathcal{A}_{xt}$. Evaluating equation (24) at time zero pins down K_0. A final issue is to guarantee that the present discounted value of utility is finite. With log utility this is indeed the case as long as all of the growth rates are bounded from above.[14]

It is worth emphasizing that condition (24) remains surprisingly tractable. In fact, it has the same functional form as the condition for GBGP in the model of Ngai and Pissarides (2007). This is surprising given the fundamental difference between their model and our model. Specifically, Ngai-Pissarides assumed that all investment is manufacturing output, whereas in our model investment is a composite of value added from the goods and services sectors. There are two reasons why the functional form of condition (24) is nonetheless the same in both models. The sectoral production functions are of the Cobb-Douglas form with the same capital-share parameter in both models, implying that in equilibrium the sectoral capital-labor ratios equal the aggregate capital-labor ratio. This implies that one can aggregate in both models, and so the GBGP condition can be written in terms of the aggregate capital-to-labor ratio. For our model, Proposition 1 above shows that this also implies that investment output can be written in terms of a sectoral pseudo-production function of the Cobb-Douglas form with the same share parameter. Given that the rental price of capital is the derivative of that sectoral pseudo-production function with respect to capital, sectoral variables enter the GBGP condition only through the effective TFP term and the sectoral composition does not enter the GBGP condition at all. This equilibrium property makes our model analytically tractable.

[14]To see this note that with log utility, the present value of discounted utility is finite as long as the growth rate of consumption is bounded. Given that $\mathcal{A}_{xt}$ grows at a constant rate, the growth rate of consumption is bounded by the difference between this growth rate and the growth of the price of consumption relative to investment. If all TFP growth rates are bounded then this rate will also be bounded.

Put somewhat differently, our analysis shows that one can study structural change separately from balanced growth. That is, the analysis of balanced growth relies sole on the pseudo production functions for aggregate consumption and investment, just as in a standard two-sector model. Conditional on establishing a balanced growth path, one can then analyze structural change. This result requires that one identifies the important role played by $\mathcal{A}_{xt}$, the *effective TFP* of the investment sector.

6 Three Insights from our Unified Approach

In this section, we present three insights that arise from our unified approach to structural change in investment and consumption. As mentioned in the introduction, these are: that technical change is endogenously investment-biased; that constant TFP growth in all sectors is inconsistent with structural change and aggregate balanced growth; and that the sector with the slowest TFP growth absorbs all resources asymptotically. Although we derived a representation of our model that makes it appear like a standard two-sector model, the first two insights are in direct contrast to the predictions of that model. In both cases the key point that explains the divergence in predictions is that the effective TFPs in our two-sector representation are in turn non-linear functions of the three primitive TFPs. The third insight relates to structural change and so has no counterpart in the standard two sector model. But our result is in contrast to models of structural change that abstract from structural change in investment.

6.1 Insight 1: Investment-Biased Technological Change

Recall the expression we derived previously for P_{ct} – the price of consumption relative to investment – as a function of primitives:

$$P_{ct} = \frac{A_{xt}\left(\omega_x A_{gt}^{\varepsilon_x - 1} + (1-\omega_x)A_{st}^{\varepsilon_x - 1}\right)^{\frac{1}{\varepsilon_x - 1}}}{\left(\omega_c A_{gt}^{\varepsilon_c - 1} + (1-\omega_c)A_{st}^{\varepsilon_c - 1}\right)^{\frac{1}{\varepsilon_c - 1}}}$$

As noted previously, in the standard two-sector model one obtains that $P_{ct} = \tilde{A}_{xt}/\tilde{A}_{ct}$, i.e., the relative price of consumption equals the TFP of the investment sector relative to the consumption sector. The change in $\tilde{A}_{xt}/\tilde{A}_{ct}$ over time is a measure of investment-biased technical change, and the previous result says that relative price movements directly reflect this bias in technical change.

One can see from equation (11) that in the special case in which $\omega_c = \omega_x$ and $\varepsilon_c = \varepsilon_x$, our model would similarly predict that $P_{ct} = A_{xt}$. Recall that since we have normalized A_{ct} to unity, this has the same form as in the standard model. But in general, our model features additional forces that impact on this relative price and one cannot say whether the additional forces lead to higher or lower growth in P_{ct}. But for now, consider the case in which $\omega_x > \omega_c$, $\varepsilon_c = \varepsilon_x$ and A_{gt} grows faster than A_{st}. We argue later that this specification is empirically reasonable. In this case, equation (11) implies that P_{ct} would increase even absent any change in A_{xt}. This effect is driven by the *endogenous* input choices for the firms producing consumption and investment. Because the investment producing firm *choose*s to have a greater ratio of goods to services, a given decrease in the relative price of goods leads to a greater decline in their minimum cost and hence a decline in their relative price. Because the literature has identified investment biased technical change with the exogenous movements in A_{xt} we refer to this additional channel in our model as *endogenous investment biased technical change*.[15] We note that this result holds along any equilibrium path and is not restricted to behavior along a GBGP.

As additional perspective on this result, recall that we also showed that the equilibrium price of consumption relative to investment is the inverse of the *effective* TFP levels in our pseudo production functions. Importantly, exogenous investment-biased technical change (i.e., A_{xt}) is only one component of this ratio.

Garcia-Santana et al. (2016) also note that the price of investment relative to consumption is influenced by the relative prices of the outputs of the underlying sectors in addition to exogenous investment specific technical change. Our analysis goes one level deeper by linking this relative price to underlying TFP levels.[16] Below, we will evaluate this mechanism empirically.

6.2 Insight 2: Sectoral TFP Growth and GBGP

In the standard model, all that is required for the existence of a GBGP is that $\tilde{A}_{xt}$ grows at a constant rate that is not too large. In particular, one can look for a GBGP assuming that both $\tilde{A}_{xt}$ and $\tilde{A}_{ct}$ grow at constant rates. Put somewhat differently, one can impose that all TFPs grow at constant but potentially different rates when looking for a GBGP. While our model has the seemingly analogous requirement that constant growth in $\mathcal{A}_{xt}$ is required, it turns out that his has a very different implication for what is allowed for the primitive TFPs. In particular, the next Proposition establishes that constant though potentially different growth rates in all TFPs are inconsistent with a GBGP that exhibits structural change.

Proposition 5 *If $\varepsilon_x \neq 1$, a necessary condition for the existence of a BGP with structural change is that at least one of the growth rates $\widehat{A}_{xt}$, $\widehat{A}_{gt}$, $\widehat{A}_{st}$ is not constant.*

[15] In the introduction, we referred to sector-specific TFP as "effective TFP".

[16] A similar insight is noted in Duernecker and Herrendorf (2015). They show that if sectoral labor input is a CES aggregator of two labor inputs that are subject to different technological progress, then sectoral differences in the relative weights of the labor inputs imply that technological progress becomes endogenously sector specific. Be that as it may, they also show that it matters for aggregate productivity growth which sector takes over, if one measures GDP as the BEA does by a Fisher index.

As previously noted, for the existence condition to be satisfied,

$$\mathcal{A}_{xt} = A_{xt}(\omega_x A_{gt}^{\varepsilon_x - 1} + (1-\omega_x) A_{st}^{\varepsilon_x - 1})^{\frac{1}{\varepsilon_x - 1}}$$

must grow at a constant rate. And a necessary condition for structural change to occur is that $\varepsilon_x \neq 1$. It then follows that if structural change occurs along a GBGP, then at most two of the three growth rates of A_{xt}, A_{gt}, and A_{st} can be constant in general. Specifically, if we were to restrict attention to constant rates of investment-specific technological change and of technological change in both the goods and services sector, the existence condition could only hold if both the goods and services sectors experienced the same rate of technological progress. But that would imply that there would not be any changes in the price of services relative to goods and hence there would not be any structural change.

While the literature on structural change and balanced growth has tended to focus on the case in which sectoral TFPs grow at constant rates, there is no natural reason to favor this case. In a one-sector model, constant growth at the aggregate level necessarily requires constant growth in aggregate TFP. But in a multi–sector model with non-linear aggregators, the fact that certain aggregates grow at (approximately) constant rates does not translate to (approximately) constant rates of growth in sectoral TFPs. In fact, the previous result shows that not only is it restrictive to only consider constant growth rates in sectoral TFPs, but that imposing this restriction will make it impossible to find a GBGP that exhibits structural change.[17] In the next section, we examine the behavior of $\mathcal{A}_{xt}$ and its various components in the data.

6.3 Insight 3: Asymptotic Behavior

Baumol (1967) was the first to note that if resources move systematically from sectors with high productivity growth to sectors with low productivity growth, then the economy will exhibit a secular decline in aggregate productivity growth. Taken to the extreme, this argument suggests that if there is one sector with low, or even zero, productivity growth, it will eventually dominate the entire economy. Importantly, Baumol's formal analysis abstracted from capital.

In contrast, Ngai and Pissarides (2007) built a model of structural change with capital and showed that it has a different asymptotic implication. Among the many sectors of their model, each of which produces a different consumption good, only one sector is assumed to produce capital. Since the capital-to-output ratio is constant along a GBGP, the capital-producing sector does not disappear, even asymptotically. Making the usual assumption that the capital-producing sector is manufacturing, and noting that in the data the manufacturing sector has stronger productivity growth than many stagnant services sectors, their model implies that while

[17] Herrendorf and Valentinyi (2015) construct a model of endogenous technological change that is consistent with this implication of the current model, in that the endogenous equilibrium growth rates of the sectoral TFPs are not constant in their model.

the sector with the slowest productivity growth will dominate consumption, it will not dominate the entire economy.

To translate this result to our setting, recall that we have only two value–added producing sectors, that is, the goods sector and the services sector. A version of their result obtains in our model if we assume that investment is produced only with goods, that goods and services are complements in the production of consumption, that TFP grows more strongly in the goods sector, and that in the limit the growth rate of goods TFP is larger than the growth rate of services TFP. While the service sector then takes over consumption asymptotically, the goods sector continues to produce investment. Since along the GBGP, investment constitutes a constant fraction of GDP, the share of the goods sector will converge to that of investment, and so the goods sector will not disappear asymptotically.

Consider instead our more general specification in which both goods and services are used to produce investment. Additionally, assume again that goods and services are complements in the production of both investment and consumption, that TFP grows more strongly in the goods sector, and that in the limit the growth rate of goods TFP is larger than the growth rate of services TFP. Under these assumptions, it remains true that investment will be a constant fraction of output, but the value added share of services will approach one for both consumption and investment. Therefore, all labor will asymptotically be in the services sector, which in our model is the sector with the slowest TFP growth. The next proposition summarizes this result:

Proposition 6 *If the parameters are such that a GBGP exists,* $\varepsilon_x, \varepsilon_c \in [0,1)$, $\widehat{A_{gt}} > \widehat{A_{st}}$, $\lim_{t\to\infty} \widehat{A_{gt}} > \lim_{t\to\infty} \widehat{A_{st}}$, *then there are two cases*

- $\omega_x = 1$: $\quad \lim\limits_{t\to\infty} \dfrac{P_{st}Y_{st}}{Y_t} = \dfrac{C}{Y}$ *(standard result);*
- $\omega_x \in [0,1)$: $\quad \lim\limits_{t\to\infty} \dfrac{P_{st}Y_{st}}{Y_t} = 1$ *(novel result).*

The proposition implies that if one takes structural change in investment into account, then the sector with the slowest productivity growth will take over out simple two-by-two-by-two economy asymptotically. This result is important when one seeks to understand the role that structural change plays for the productivity slowdown.[18]

7 Empirical Analysis

In this section we complement the preceding theoretical analysis of the model of Section 3 with an empirical assessment of two questions. First, to what extent are the CES aggregators in our

[18] In related work, Duernecker et al. (2017b) point out that in a more disaggregated setting, the sector with the slowest productivity growth may not take over the economy asymptotically, because services with fast and slow productivity growth turn out to be substitutes, instead of complements.

Table 1: Preference and technology parameters

ω_x	ω_c	ε_x	ε_c
0.52	0.19	0.00	0.00

formulation able to capture the secular trends in the post WWII US data? Second, to the extent that we observe something that looks like approximate balanced growth in the data, what is the nature of the technological change that accounts for this. Put somewhat differently, what is the behavior of the various terms in the theoretical condition necessary for generalized balanced growth?

7.1 Fit of the CES Aggregators

We begin with the first question. Our analysis focuses on the expressions for goods and services expenditure shares by final expenditure sector that we previously derived, which we repeat here for convenience:

$$\frac{P_{gt}X_{gt}}{P_{st}X_{st}} = \frac{\omega_x}{1-\omega_x}\left(\frac{P_{gt}}{P_{st}}\right)^{1-\varepsilon_x} \quad \text{and} \quad \frac{P_{gt}C_{gt}}{P_{st}C_{st}} = \frac{\omega_c}{1-\omega_c}\left(\frac{P_{gt}}{P_{st}}\right)^{1-\varepsilon_c}.$$

To assess whether the assumed CES structure is empirically reasonable we ask whether there are values for the share and elasticity parameters (the ω_j and ε_j) such that when taking relative prices as given by the data, we are able to capture the key secular changes in expenditure shares. To implement this we use the data as presented in Section 2 and calibrate the share and elasticity parameters to minimize the sum of squared deviations between the implied expenditure shares and the expenditure shares in the data. This is effectively the same procedure that we used in Herrendorf et al. (2013) when estimating a CES aggregator for consumption with three components–agriculture, non-agricultural goods, and services. The implied values for the current exercise are presented in Table 1.

As expected, the relative weight of goods is estimated to be higher in investment than in consumption; compare Figures 1. The striking result is that *both* elasticity parameters turn out to be zero, implying that both aggregators are Leontief. While the result that $\varepsilon_c = 0$ is somewhat expected given the results obtained earlier in Herrendorf et al. (2013), it is perhaps surprising that the aggregator in the investment sector features a similarly low degree of substitutability.

To understand why Leontief aggregators provide the best fit for the composition of final expenditures in the postwar US we consider the ratio of the *real* inputs of goods to services in each sector. With CES aggregators and inputs being complements, real ratios move in the

Figure 3: Real composition of final expenditures

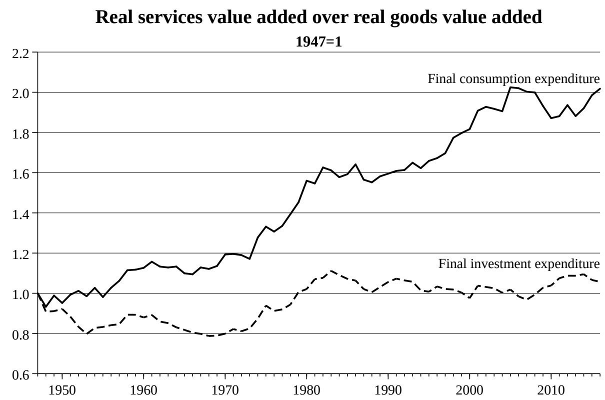

opposite direction of expenditure shares. That is, with an elasticity of substitution less than unity, an increase in the relative price of services implies that the expenditure share of services increases, but that the ratio of real services to real goods decreases. Intuitively, the reallocation of relative real quantities serves to the dampen the change in expenditure shares, but it does not overturn the initial impact of the change in prices. In the extreme case of a Leontief aggregator, there is a constant ratio for real quantities.

With this in mind Figure 3 displays the time series for ratios of real inputs. The figure shows that for investment, the ratio of real services to real goods inputs increases modestly. Given that the relative price of services to goods increases quite dramatically over the period of investigation, this points to a Leontief specification, which keeps the real composition unchanged. This suggests that the Leontief specification actually seems to be the natural specification for modeling structural change in investment.

Moving to consumption, we observe that the ratio of real services to real goods inputs increases significantly at the same time that the relative price of services to goods also increases significantly. Given our earlier discussion, a CES aggregator is unable to capture this, because the most extreme case is the Leontief specification, in which case the real composition remains unchanged in the face of a large change in relative prices. In all other specifications with an elasticity of substitution less than unity, the real quantity of services would in fact *decrease*. This suggests that, again, the Leontief specification will do the best possible job, though it will still miss an important part of the changes in the real composition of inputs into consumption. However, as we noted earlier, the latter result is to be expected since we have abstracted from income effects by virtue of considering a homothetic aggregator. Given that our earlier work that focused solely on consumption did find a role for income effects, including them would add a force leading to an increase in the ratio of real services to real goods in the production of consumption.

Figure 4: Expenditure shares – data versus model

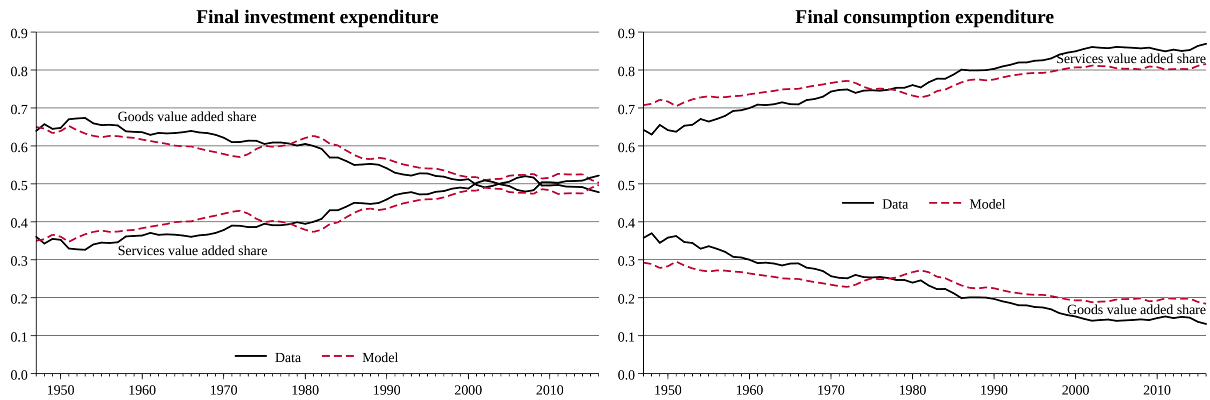

Figures 4 show the fitted values for the expenditure shares along with their corresponding values in the data. As suggested by the previous discussion, we find that the fitted values track the secular change within the investment sector very well, whereas the fitted values for consumption explain only about half of the change in expenditure shares within consumption.[19] Additionally, recall that we have aggregated agriculture and manufacturing into a single goods sector and that nonhomotheticities are typically found to be of particular importance for the decline of agriculture. Given that our primary focus in this paper was on structural change within investment, and that we abstracted from non-homotheticities within consumption in order to facilitate transparency, we are not overly concerned about missing an important part of the structural change in consumption here. Importantly, the figures show that our CES aggregator does a very good job of capturing the secular change within investment in the US over this time period.

7.2 Behavior of $\mathcal{A}_{xt}$

We now turn our attention to the necessary condition for balanced growth – that the effective TFP in the investment sector, $\mathcal{A}_{xt}$, grow at a constant rate – from an empirical perspective. To do this we use the estimated values for ω_x from the previous exercise and compute series for the $A_{jt}'s$ using standard growth accounting methods as in Solow (1957). Specifically, to solve for A_{jt} ($j \in \{g, s\}$), we set $A_{j0} = 1$ and calculate the growth rates by using the information provided by WORLD KLEMS on nominal value added, capital services, labor services in efficiency unit, and factor shares:

$$\widehat{P_{jt}Y_{jt}} = \widehat{A}_{jt} + \frac{R_t K_{jt}}{P_{jt}Y_{jt}}\widehat{K}_{jt} + \frac{W_t L_{jt}}{P_{jt}Y_{jt}}\widehat{L}_{jt}.$$

[19]We note that adding net exports to consumption tends to worsen the fit of the model.

Figure 5: Exogenous TFPs

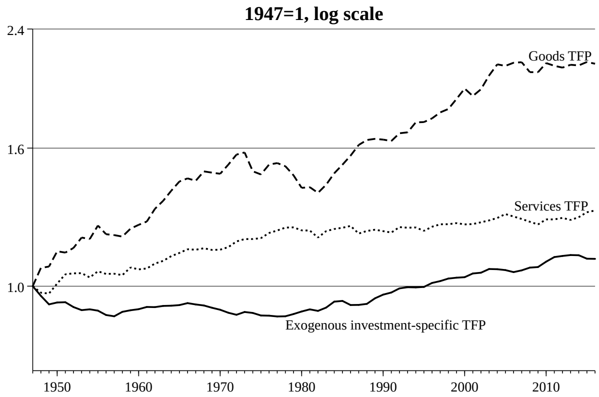

Since all terms except for $\widehat{A}_{jt}$ are observable, we can solve for $\widehat{A}_{jt}$ and, given the normalization $A_{j0} = 1$, construct A_{jt}. To estimate TFP growth in the investment sector we use:[20]

$$\widehat{X}_t = \widehat{A}_{xt} + \frac{P_{gt}X_{gt}}{X_t}\widehat{X}_{gt} + \frac{P_{st}X_{st}}{X_t}\widehat{X}_{st}$$

to compute the implied growth rate series for $\widehat{A}_{xt}$. Using this series together with the series for $\widehat{A}_{gt}$ and $\widehat{A}_{st}$ and our calibrated parameter values, we can uncover the series for $\mathcal{A}_{xt}$ from the expression that defines $\mathcal{A}_{xt}$:

$$\underbrace{\mathcal{A}_{xt}}_{\text{investment-specific TFP}} \equiv \underbrace{A_{xt}}_{\text{first component}} \underbrace{\left(\omega_x A_{gt}^{\varepsilon_x-1} + (1-\omega_x)A_{st}^{\varepsilon_x-1}\right)^{\frac{1}{\varepsilon_x-1}}}_{\text{second component}}.$$

Note that we impose our functional forms only when we compute $\widehat{A}_{xt}$, but not when we compute $\widehat{A}_{gt}$, $\widehat{A}_{st}$, and $\widehat{\mathcal{A}}_{xt}$.

Figure 5 shows the implied series for the level of each TFP term. The graph has the interesting implication that there is little "pure" investment-specific technological change arising from the first component of $\mathcal{A}_{xt}$. Instead, higher TFP growth in the investment sector is driven by the second component. As we explained in Insight 1, technological progress is endogenously investment biased in our model because the effective TFP puts a higher weight on goods in investment than in consumption, and goods have higher TFP growth than services.

Given our estimates of ω_{xj} for $j \in \{g, s\}$ and ε_x, plus our estimated series for the A_{jt} for $j \in \{x, g, s\}$, it is straightforward to compute the value of $\mathcal{A}_{xt}$. Recall that constant growth in this quantity is necessary in order to have a GBGP. The left panel of Figure 6 plots the log of this value over time. The figure shows that this term indeed grows on average at a constant rate

[20]The continuous–time divisia index was approximated with the Törnqvist index in the data.

Figure 6: Investment TFPs

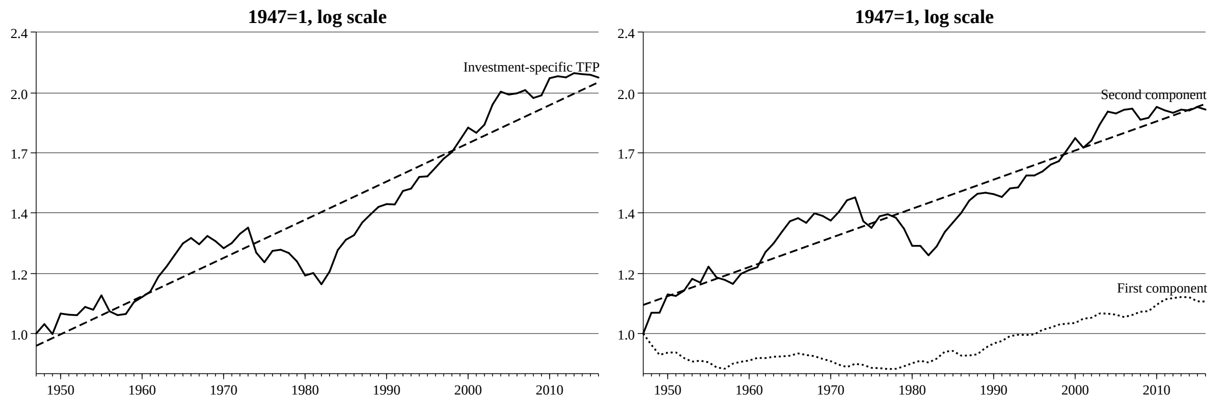

over the postwar period.

We emphasize two points. First, the construction of this figure did not directly use the information that the average growth of chained GDP in the US data is approximately constant. Second, although $\mathcal{A}_{xt}$ does grow at an approximately constant rate over the time period being studied, none of the individual TFP terms exhibits approximate constant growth. In particular, while the growth in A_{st} declines over time, the growth in A_{xt} actually increases over time. Note moreover, that the increase in A_{xt} is rather small.

8 Disaggregating Investment

To this point we have assumed a single consumption good and a single investment good in order to best illustrate the main insights that derive from considering structural change within the investment sector. In this section we consider a further disaggregation within the investment sector to assess the extent to which the previously documented patterns reflect compositional changes across categories of investment versus changes within specific categories. To this end we disaggregate total investment into three categories: structures, equipment, and intellectual property product (IPP henceforth). The BEA counts three types of IPP investments: software, research and development, and artistic products.

Figure 7 shows the evolution of the shares of total investment accounted for by each of these three categories. The dominant secular pattern is a steady decline in the share of investment accounted for by structures, totalling roughly fifteen percentage points since 1947. There is also a steady downward trend for the share of total investment that is accounted for by equipment, though this series also has some low frequency fluctuations. Offsetting these two downward trends is the steady increase in the share of investment accounted for by IPP. While at the beginning of the sample structures accounts for more than half of total investment and IPP

Figure 7: Expenditure Shares of Subcategories of Investment

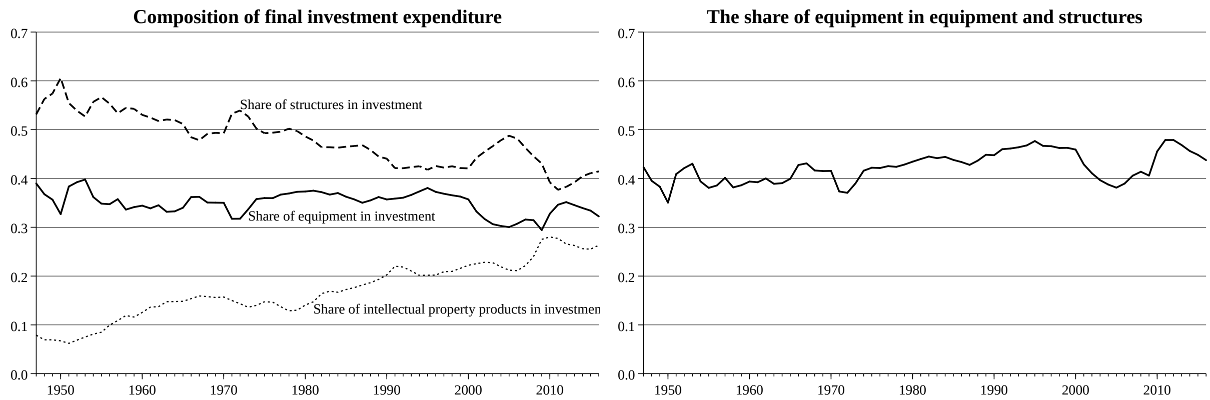

accounts for less than ten percent, at the end of the sample structures accounts for roughly 40% of the total with each of the other two categories accounting for roughly 30% each. For future purposes, we show in the right panel of Figure 7 that if we focus purely on structures and equipment, the share of this total accounted for by structures is relatively constant over the sample period.

The three panels in Figure 8 display the extent of structural change within each of the three investment categories. Interestingly, we see that structural change occurs within both equipment and IPP, but is absent in structures. Additionally, note that the initial share of services is much higher for IPP than for the other two categories. As of the end of the sample period, services accounts for more than half of final equipment expenditure and for over 90% of final IPP expenditure. To first approximation, the series for equipment mimics the series for aggregate investment that we displayed earlier in the paper.

The left panel of Figure 9 displays the evolution of prices for structures and IPP relative to equipment. The key message from this figure is that the price of structures has increased dramatically relative to the prices of the two other categories. While the price of IPP has also increased relative to equipment, the increase is quite modest compared to the increase in the relative price of structures.

Combining data on relative shares and relative prices we can assess the secular trends in the quantity of investment in each of the three categories. This information is presented in the right panel of Figure 9, where we have normalized the initial level for each category equal to unity. As the figure shows, IPP and equipment are both increasingly important in terms of accounting for the quantity of investment.

We draw two main messages from the above empirical patterns. First, the phenomenon of structural change within the investment sector is strongly present in both equipment and IPP. That is, structural change within investment is not purely a phenomenon of changing composi-

Figure 8: Sector Shares in the Subcategories of Investment

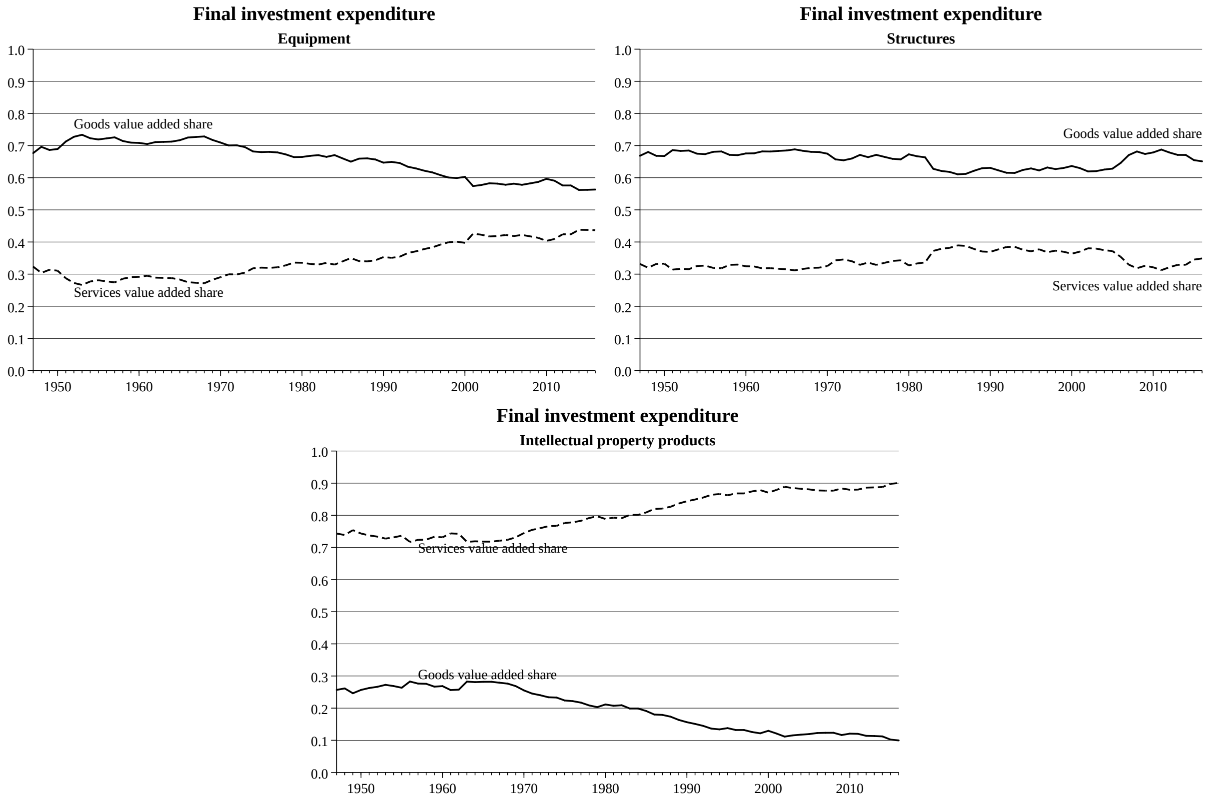

Figure 9: Relative Prices of the Subcategories of Investment

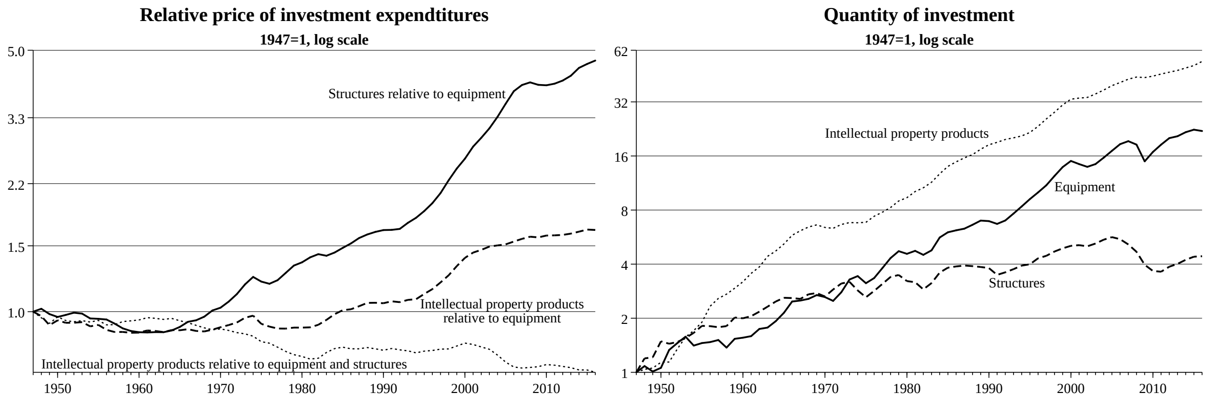

tion within investment. Second, there are important secular trends in both the composition of investment across categories and the relative prices of the three categories of investment.

Our previous analysis abstracted from compositional changes within the investment sector. Given that the extent of structural change varies across the three categories we think the above empirical facts do suggest that it is of interest to extend our analysis to consider three different categories of investment. We leave a full analysis of this to future work, because disaggregating investment destroys the analytical tractability of our simple two-by-two-by-two structure and implies that one needs to simulate the model. We leave this for future work, because our point in this paper is to characterize in the simplest, most intuitive way the key forces that are at work with structural change in investment. In what follows, we therefore only map out a formulation of the more general model that is motivated by the above empirical patterns.

As a first step in extending the previous framework, we note that it is straightforward to assume that there are now four categories of final expenditure: consumption and the three categories of investment. Consistent with our previous framework, we can express the production of each of these four categories of final expenditure as CES functions of value added from the goods and services sectors. In particular, we would have:

$$X_{it} = A_{xit}\left(\omega_{xi}^{\frac{1}{\varepsilon_{xi}}} X_{git}^{\frac{\varepsilon_{xi}-1}{\varepsilon_{xi}}} + (1-\omega_{xi})^{\frac{1}{\varepsilon_{xi}}} X_{sit}^{\frac{\varepsilon_{xi}-1}{\varepsilon_{xi}}}\right)^{\frac{\varepsilon_{xi}}{\varepsilon_{xi}-1}}$$

where $i \in \{1, 2, 3\}$ and the three indices reflect structures, equipment and IPP respectively. Note that we allow the weights and the elasticity of substitution to vary across the three categories, as well as for each category to have its own level of technical progress. As noted previously, in the face of large changes in the price of services relative to goods, the expenditure share for services has increased in both equipment and IPP, but has remained constant in structures. The same logic that we applied previously would lead us to conclude that the elasticity of substitution is less than unity for both equipment and IPP but is equal to unity for structures. That is, the aggregator for structures is Cobb-Douglas. Similarly, the weight on services will be greatest for IPP and smallest for structures. Finally, given these parameters, the behavior of relative prices of the three categories would provide information on the category specific rate of technical progress.

The second step in the extension needs to specify how the three types of capital enter into the production function. The existing literature has considered this for the case in which there are two categories of capital – equipment and structures, and used different formulations. Greenwood et al. (1997) effectively used a Cobb-Douglas aggregator to aggregate the two types of capital into a single capital good. In contrast, Krusell et al. (2001) allowed for two different labor inputs and assumed that equipment and structures had differential substitutability with the two different types of labor, and so do not aggregate the two capitals into a single stock.

For present purposes we continue to assume a single type of labor and look to aggregate the three capital categories into a single stock. The previously presented facts suggest that a single elasticity of substitution between the three categories will not be sufficient. To see why note that both IPP and equipment have decreasing prices relative to structures, but that whereas expenditure on equipment is roughly constant relative to structures, expenditure on IPP increases quite dramatically. In view of this there are of course multiple ways to nest the aggregation of the three different capitals. Here we present one simple form that seems intuitive given the empirical patterns noted above. In particular, the fact that the share of equipment in the aggregate of equipment and structures is roughly constant in the face of large changes in the relative price of the two, suggests a Cobb-Douglas aggregator of the two. This aggregate would then be combined with IPP using a CES aggregator. The previous evidence suggests that this aggregator would have an elasticity of substitution greater than unity. Specifically, the IPP share in total investment rises while the price of IPP relative to the structures/equipment composite falls; compare the left panels of Figures 7 and 9.

The above discussion suggests an empirically reasonable specification for an analysis that wants to consider a disaggregation of investment into three categories. Given these functional forms, it is not possible to derive analytic expressions to characterize the aggregate behavior and so any further analysis will require numerical simulations. We leave this analysis for future work.

As final point we note one other issue that the disaggregation of investment into structures, equipment, and IPP naturally brings up. This issue is that IPP capital is part of intangible capital, which has some nonrival components because the underlying ideas can often be used by everyone.[21] To avoid confusion of what this means in the context of IPP capital, it is crucial to make the distinction between the concepts of nonrival and excludable. Nonrivalry is a feature of nature whereas excludability is a feature of the institutions chosen by society. Nonrivalry means that the use of an idea by one person does not diminish its usefulness for other persons. Excludability means that property rights allow the legal restriction of use to one person. Irrespective of whether IPP capital is rival or nonrival, it is excludable. That is, the BEA counts as IPP capital only those investments that result in well defined property rights and have observable market prices. For software, property rights are established by copyrights and by tailoring the software to the particular user; for research and development, property rights are established by patents; and for artistic products property rights are established by copyrights. The fact that IPP capital has well defined property rights implies that, in the context of the standard growth model, one may treat its stock in the same way as we treat the stock of any other type of capital.

[21] For recent attempts to model intangible capital, see McGrattan and Prescott (2009,2010) and Xi (2016).

9 Conclusion

In this paper, we have proposed a new framework for studying structural change and growth that treats consumption and investment in a symmetric manner. In particular, we have modelled production at the level of sector value added and have treated both final consumption and final investment as aggregates of the underlying sectoral value added. We have studied a simple form of this framework with two final demand sectors – consumption and investment – and two underlying sectors producing value added – goods and services. We have focused on this simple framework to best able to illustrate its distinctive features. First, for empirically plausible parameter values technological change is endogenously investment specific. Second, constant growth in all sectoral TFPs is generically inconsistent with aggregate balanced growth. Third, the sector with the lowest productivity growth asymptotically dominates the entire economy. Lastly, we have shown that a version of the model with a CES aggregator in the investment sector can account for the salient trends in the value-added composition of final investment.

We believe that richer versions of this model will prove useful in refining our view of the nature of structural change and the forces that drive it. In particular, we think it will be of interest to pursue more fully a formulation that separates investment into the three components of structures, equipment and IPP. We believe this framework will help us better isolate the underlying sources of TFP growth.

Lastly, in this paper, we have abstracted from non-homotheticities in the aggregator of value added from different sectors to final consumption. Our empirical analysis has confirmed that these non-homotheticities play a quantitatively important role. It is therefore of interest to explore to what extent one can integrate them in the present framework without losing tractability. We plan to turn to this task next.

References

Acemoglu, Daron and Veronica Guerrieri, "Capital Deepening and Non–Balanced Economic Growth," *Journal of Political Economy*, 2008, *116*, 467–498.

Alvarez-Cuadrado, Francisco, Ngo Van Long, and Markus Poschke, "Capital–Labor Substitution, Structural Change, and Growth," *Theoretical Economics*, 2017, *12*, 1229–1266.

Baumol, William J., "Macroeconomics of Unbalanced Growth: The Anatomy of the Urban Crisis," *American Economic Review*, 1967, *57*, 415–426.

Boppart, Timo, "Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non–Gorman Preferences," *Econometrica*, 2014, *82*, 2167–2196.

Comin, Diego, Martí Mestieri, and Danial Lashkari, "Structural Transformations with Long–run Income and Price Effects," Manuscript, Northwestern University 2015.

Duernecker, Georg and Berthold Herrendorf, "Occupations, Barriers to Entry and Hours Worked in the U.S. and Germany," Manuscript, University of Mannheim and Arizona State University 2015.

——, ——, and Ákos Valentinyi, "Quantity Measurement and Balanced Growth in Multi–Sector Models," Manuscript 2017.

——, ——, and ——, "Structural Change within the Service Sector and the Future of Baumol Disease," Manuscript 2017.

Echevarria, Cristina, "Changes in Sectoral Composition Associated with Economic Growth," *International Economic Review*, 1997, *38*, 431–452.

Garcia-Santana, Manuel, Josep Pijoan-Mas, and Lucciano Villacorta, "Investment Demand and Structural Change," Manuscript 2016.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, "Long-run Implication of Investment-Specific Technological Change," *American Economic Review*, 1997, *87*, 342–362.

Herrendorf, Berthold and Ákos Valentinyi, "Endogenous Sector–Biased Technological Change and Industrial Policy," Discussion Paper 10869, CEPR, London 2015.

——, Christopher Herrington, and Ákos Valentinyi, "Sectoral Technology and Structural Transformation," *American Economic Journal: Macroeconomics*, 2015, *7*, 1–31.

——, Richard Rogerson, and Ákos Valentinyi, "Two Perspectives on Preferences and Structural Transformation," *American Economic Review*, 2013, *103*, 2752–2789.

——, ——, and ——, "Growth and Structural Transformation," in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth*, Vol. 2, Elsevier, 2014, pp. 855–941.

Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie, "Beyond Balanced Growth," *Review of Economic Studies*, 2001, *68*, 869–882.

Krusell, Per, Lee E. Ohanian, José-Víctor Ríos-Rull, and Giovanni L. Violante, "Capital-Skill Complementarities and Inequality: a Macroeconomic Analysis," *Econometrica*, 2001, *68*, 1029–1054.

McGrattan, Ellen R. and Edward C. Prescott, "Openness, Technology Capital and Development," *Journal of Economic Theory*, 2009, *144*, 2454–2476.

—— **and** ——, "Unmeasured Investment and the Puzzling US Boom in the 1990s," *American Economic Journal: Macroeconomics*, 2010, *2*, 88–122.

Ngai, L. Rachel and Christopher A. Pissarides, "Structural Change in a Multisector Model of Growth," *American Economic Review*, 2007, *97*, 429–443.

Solow, Robert, "Technical Change and the Aggregate Production Function," *Review of Economics and Statistics*, 1957, *38*, 312–320.

Xi, Xican, "Multi-Establishment Firms, Misallocation and Productivity," Manuscript, Fudan University 2016.

We emphasize that characterizing a GBGP is more challenging in our framework than usual, because we have structural change in both the consumption sector and the investment sector. In contrast, existing models of structural change typically abstract from structural change within the investment sector. Doing this has the advantage that it "anchors" the economy and makes it fairly straightforward to obtain a constant rental price of capital by imposing that TFP growth in investment, which is a primitive of the existing models, is constant. Things are rather different in our framework in which TFP in investment is an endogenous, non–linear aggregator of the other TFPs.

Appendix

Proof of Proposition 1

We show the claim for X_t only. The proof for C_t follows the exact same steps. We start by rewriting the production function of X_t:

$$X_t = A_{xt}\left(\omega_x^{\frac{1}{\varepsilon_x}} + (1-\omega_x)^{\frac{1}{\varepsilon_x}}\left(\frac{X_{st}}{X_{gt}}\right)^{\frac{\varepsilon_x-1}{\varepsilon_x}}\right)^{\frac{\varepsilon_x}{\varepsilon_x-1}} X_{gt},$$

$$X_t = A_{xt}\left(\omega_x^{\frac{1}{\varepsilon_x}}\left(\frac{X_{gt}}{X_{st}}\right)^{\frac{\varepsilon_x-1}{\varepsilon_x}} + (1-\omega_x)^{\frac{1}{\varepsilon_x}}\right)^{\frac{\varepsilon_x}{\varepsilon_x-1}} X_{st}.$$

Equation (20) implies that:

$$\frac{X_{st}}{X_{gt}} = \frac{1-\omega_x}{\omega_x}\left(\frac{A_{st}}{A_{gt}}\right)^{\varepsilon_x} \quad \text{and} \quad \frac{X_{gt}}{X_{st}} = \frac{\omega_x}{1-\omega_x}\left(\frac{A_{gt}}{A_{st}}\right)^{\varepsilon_x}.$$

Combining the last two sets of equations, we obtain:

$$
\begin{aligned}
X_t &= A_{xt}\left(\omega_x^{\frac{1}{\varepsilon_x}} + \frac{1-\omega_x}{\omega_x^{\frac{\varepsilon_x-1}{\varepsilon_x}}}\left(\frac{A_{st}}{A_{gt}}\right)^{\varepsilon_x-1}\right)^{\frac{\varepsilon_x}{\varepsilon_x-1}} X_{gt},\\
X_t &= A_{xt}\left(\frac{\omega_x}{(1-\omega_x)^{\frac{\varepsilon_x-1}{\varepsilon_x}}}\left(\frac{A_{gt}}{A_{st}}\right)^{\varepsilon_x-1} + (1-\omega_x)^{\frac{1}{\varepsilon_x}}\right)^{\frac{\varepsilon_x}{\varepsilon_x-1}} X_{st}.
\end{aligned}
$$

Rearranging gives:

$$
\begin{aligned}
\omega_x A_{gt}^{\varepsilon_x-1} X_t &= A_{xt}\left(\omega_x A_{gt}^{\varepsilon_x-1} + (1-\omega_x)A_{st}^{\varepsilon_x-1}\right)^{\frac{\varepsilon_x}{\varepsilon_x-1}} \frac{X_{gt}}{A_{gt}},\\
(1-\omega_x) A_{st}^{\varepsilon_x-1} X_t &= A_{xt}\left(\omega_x A_{gt}^{\varepsilon_x-1} + (1-\omega_x)A_{st}^{\varepsilon_x-1}\right)^{\frac{\varepsilon_x}{\varepsilon_x-1}} \frac{X_{st}}{A_{st}}.
\end{aligned}
$$

Adding these two equations, we obtain:

$$
\left(\omega_x A_{gt}^{\varepsilon_x-1} + (1-\omega_x)A_{st}^{\varepsilon_x-1}\right) X_t = A_{xt}\left(\omega_x A_{gt}^{\varepsilon_x-1} + (1-\omega_x)A_{st}^{\varepsilon_x-1}\right)^{1+\frac{1}{\varepsilon_x-1}} \left(\frac{X_{gt}}{A_{gt}} + \frac{X_{st}}{A_{st}}\right)
$$

Dividing through the first term on the left–hand side and using the definitions of $\mathcal{A}_{xt}$, K_{xt}, and L_{xt} proves the claim. **QED**

Proof of Proposition 3

Note first that if R_t is constant, the household's Euler equation (22) implies that E_t must grow at a constant rate. Denoting this rate by γ, we have that γ satisfies:

$$
\gamma = R - \delta - \rho.
$$

Second, because both the goods and services sectors have Cobb-Douglas production functions with the same capital share, total payments to capital will be a fraction θ of total output: $RK_t = \theta Y_t$. A constant value for R thus also implies that Y_t/K_t is constant. Using this expression to substitute for R in the Euler equation gives:

$$
\frac{Y_t}{K_t} = \frac{\gamma + \delta + \rho}{\theta}. \tag{25}
$$

The fact that Y_t and K_t grow at the same rate does not necessarily imply that they grow at constant rates. But we next show that along a GBGP it is indeed the case that K_t grows at a constant rate, and furthermore that this rate is γ. Using $\dot{K}_t = X_t - \delta K_t$ with the resource constraint $X_t = Y_t - E_t$ gives:

$$\widehat{K_t} = \frac{Y_t}{K_t} - \delta - \frac{E_t}{K_t}.$$

Substituting for Y_t/K_t using equation (25) gives:

$$\widehat{K_t} = \frac{\gamma + \delta + \rho}{\theta} - \delta - \frac{E_t}{K_t}. \tag{26}$$

To show that $\widehat{K_t}$ is constant and equal to γ, we argue by way of contradiction. First, assume that at some time T $\widehat{K_t} > \gamma$. Since E_t grows at rate γ this implies that E_t/K_t is decreasing. Equation (26) then implies that $\widehat{K_t}$ is increasing, so that $\widehat{K_t}$ will exceed γ for all t greater than T. It follows that the limit of E_t/K_t must be zero and hence that:

$$\lim_{t\to\infty} \widehat{K_t} = \frac{\gamma + \delta + \rho}{\theta} - \delta = \frac{\gamma + \rho}{\theta} + \frac{1-\theta}{\theta}\delta > \gamma + \rho.$$

But this implies that the transversality condition, which requires that $\widehat{K_t} < \gamma + \rho$, is violated.

Suppose alternatively that at some time T we have $\widehat{K_t} < \gamma$. Arguing as above, E_t/K_t is now increasing, implying that $\widehat{K_t}$ must be decreasing, so that $\widehat{K_t}$ will be less than γ for all $t > T$. This implies that E_t/K_t will tend to infinity and hence that:

$$\lim_{t\to\infty} \widehat{K_t} = -\infty.$$

It follows that the growth rate of K_t is negative and bounded away from zero beyond some finite date, implying that K_t must become negative in finite time. This violates feasibility.

We conclude that K_t grows at the constant rate γ. Since Y_t grows at the same rate as K_t it follows that Y_t also grows at the same constant rate γ. It remains to show that X_t grows at rate γ. Combining $\dot{K}_t = X_t - \delta K_t$ with the fact that K_t grows at the constant rate γ gives:

$$\frac{X_t}{K_t} = \gamma + \delta,$$

so that X/K is constant. It follows that it must also be that $\widehat{X_t} = \gamma$. **QED**