Review of Economic Studies (2020) **10**, 1–37 0034-6527/20/00000001$02.00

Structural Change in Investment and Consumption – A Unified Analysis

BERTHOLD HERRENDORF

Arizona State University

RICHARD ROGERSON

Princeton University and NBER

ÁKOS VALENTINYI

University of Manchester, CERS–HAS, and CEPR

First version received July 2018; final version accepted January 2020 (*Eds.*)

The structural-change literature typically assumes that investment is produced in manufacturing. We establish that this assumption is counterfactual: in the postwar US, the share of services value added in investment expenditure has been steadilly growing. We develop a new model that features structural change in investment and consumption, characterize its equilibrium properties, and provide empirical support for it. We establish that modelling structural change in investment leads to three novel insights: constant TFP growth in all sectors is inconsistent with the existence of aggregate balanced growth with structural change; the sector with the slowest TFP growth absorbs all resources asymptotically; technical change is endogenously investment-biased.

Keywords: aggregate balanced growth, investment, investment-biased technical change, structural change
JEL Code: O11; O14.

1. INTRODUCTION

A large, recent literature has proposed extensions of the one-sector growth model 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 in investment, assuming 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 is of interest to assess the appropriateness of this assumption. In this paper, we show that the standard assumption is strongly counterfactual and affects the model's predictions in crucial ways. In other words, abstracting from structural change in investment is neither empirically plausible nor theoretically innocuous.

Our analysis begins by decomposing final investment expenditure in the post World War II US into the value added shares coming from the goods sector and the services sector. This is the analogue of the decomposition of final consumption expenditure in

1. This practice started with the early model 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.

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, the standard assumption is squarely at odds with the data; the share of services value added in investment expenditure is large, and in fact it now exceeds the share of goods value-added in investment expenditure. Second, and more importantly, there is structural change in investment expenditure, and it features the same qualitative patterns as structural change in consumption expenditure: the expenditure share of services value added 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 balanced growth and structural change in both of the final expenditure categories: investment and consumption. Our framework can accommodate various levels of disaggregation, but to best highlight the key economic interactions, our analysis focuses on the simplest two-by-two-by-two structure. This structure features two final expenditure categories (investment and consumption), two underlying sectors that produce value added (goods and services), and two factors for producing value added (capital and labor). Goods and services value added are each produced by combining capital and labor according to Cobb-Douglas production functions that are assumed to have the same capital intensity but are allowed to experience different rates of TFP growth. Investment is obtained by combining value added from the goods and services sectors, and this production function also experiences its own rate of TFP growth, which we label as exogenous investment-specific technical change. Consumption of goods and services are implicitly aggregated within the utility function of the representative household. While the aggregator for investment is assumed to be of the constant elasticity of substitution functional form (CES henceforth), our specification for consumption is consistent with two popular utility aggregators: a log-CES instantaneous utility and a price-independent-generalized-linear indirect utility (PIGL henceforth). While the former is homothetic, the latter is not.

Having developed our new model, we proceed to examine its equilibrium properties. We show that we may analyse the existence of an aggregate balanced growth path (ABGP henceforth) and the occurrence of structural change in two separate steps.[2] Key to this result is establishing that, in equilibrium, the output of the investment sector can be represented as a Cobb-Douglas production function in capital and labor with an *effective* TFP that is given by the product of exogenous investment-biased technical change and a CES aggregator of the TFPs of goods and services production. In other words, in equilibrium, our new model reduces to the model we studied previously in Herrendorf et al. (2014), which had multiple consumption sectors, one investment sector, and sectoral Cobb-Douglas production functions with the same capital intensity but different TFPs. The novelty of our current framework is that the investment sector of its reduced-form has a TFP that depends endogenously on the composition of investment inputs, which in turn depends on the TFPs of goods and services production.

Accordingly, we first provide necessary and sufficient conditions for an ABGP to exist in our model. We then establish that standard forces dictate structural change in investment and consumption. In particular, increases in the relative price of services to goods combined with a less than unitary elasticity of substitution lead to increases in the expenditure shares of services in both investment and consumption. In addition, if nonhomotheticities in the utility function imply that the income elasticity of services

2. We use ABGP in the sense introduced by Ngai and Pissarides (2007), that is, *aggregate* variables grow at constant rates including zero while *sectoral* variables may not grow at constant rates.

allocated to consumption is larger than one, then growth in aggregate income also leads to an increase in the expenditure share of services in the consumption sector. Taken together, both forces lead to a reallocation from goods to services within both final expenditure sectors. We call this the *intensive margin* of structural change. Since investment and consumption have different compositions of goods and services, aggregate structural change could also occur along the *extensive margin* due to changes in the mix of investment and consumption in aggregate final expenditure. We show that along the ABGP of our model no adjustment happens along the extensive margin and all structural change comes from the intensive margin.

The CES aggregator in the investment sector of our framework is admittedly somewhat specialized, and so there may be concern about how well it will perform quantitatively. We establish that the CES aggregator in investment captures well the salient features of structural change in the US investment sector during the post World War II period. We also establish that a log-CES specification in consumption captures part of the rising expenditure share for services but fails to account for the entire increase. In contrast, the PIGL specification captures well all of the rising expenditure share for services. Given the work of Boppart (2014), the last finding is not entirely surprising, but we note that the two exercises are not identical; our exercise matches value added shares from NIPA data whereas his matched final expenditure shares from household expenditure data. As emphasized in Herrendorf et al. (2013), switching from final expenditure shares to value added shares affects both the expenditure shares and the relative price series. An additional complication in comparing the two results is that private household expenditure data does not cover some large categories of education and health care spending.

We also carry out an empirical assessment of our new theoretical condition for the existence of an ABGP, which places a non-linear restriction on the evolution of the sectoral TFPs in our model. We calculate 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 some 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 establish that modelling structural change in the investment sector leads to three important, novel insights. First, in our model, constant (but possibly different) growth in each of the three primitive TFP terms is inconsistent with aggregate balanced growth and structural change in investment and consumption occurring jointly. Second, the sector with the slowest TFP growth asymptotically absorbs all resources in the investment sector as it does in the consumption sector. Third, technical change becomes *endogenously* investment specific, in that part of the change in the effective TFP in the investment sector results from the endogenous decisions of firms about the composition of investment inputs. We show that this endogenous component is quantitatively more important than the exogenous component.

An outline of the paper follows. In the next section, we discuss the most closely related literature. In Section 3, we present the key facts from the US time series data since World War II. Sections 4 and 5 present our model and derive key properties of the competitive equilibrium that hold irrespective of the functional form of the instantaneous utility. In Section 6, we establish conditions for the existence of an aggregate balanced growth path with structural change in both investment and consumption. Section 7 empirically examines key aspects of our new model. In Sections 8 and 9, we derive the second and third insight that result from studying structural change in investment. Section 10 concludes and an Appendix contains the proofs of our theoretical results.

2. RELATED LITERATURE

In this section, we discuss the closely related papers by Garcia-Santana et al. (2018) and Acemoglu and Guerrieri (2008), which also consider structural change in investment. Despite this similarity, we will point out the crucial differences between our work and theirs.[3]

Garcia-Santana et al. (2018) 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, their focus is very different from ours. We analytically establish the existence and properties of an aggregate balanced growth path with structural change in a general equilibrium setting. In our analysis, structural change arises entirely from the *intensive* margin, because, as we show, along an ABGP the extensive margin is not operational. In contrast, they study transition dynamics and establish empirically that the *extensive* margin is an important driver of structural change away from the balanced growth path. Specifically, they show that movements in the investment-to-GDP ratio, as occur during an investment boom, produce hump-shaped movements in the value added share of the manufacturing sector for a large sample of countries.

Acemoglu and Guerrieri (2008) develop a model in which investment and consumption feature structural change because both have the *same* sectoral composition. In contrast, the evidence we present suggests that the sectoral composition of investment and consumption differ considerably in the postwar US, in that the goods share was considerably higher in investment than in consumption expenditure. We capture this feature of the data and show that it has important implications. Also, their framework implies a constant relative price between consumption and investment, a feature that is also strongly at odds with the data. Our framework not only allows us to match the behaviour of the price of consumption relative to investment, but also provides novel insights into its underlying determinants.

There are two other important differences between Acemoglu and Guerrieri (2008) and our work. First, they assume that the instantaneous utility is homothetic in goods and services value added whereas we also study a non-homothetic instantaneous utility. Boppart (2014) and Comin et al. (2018) argue that non-homothetic instantaneous utility is essential for matching structural change in consumption. Second, they focus on the implications of differences in the capital intensities of goods and services production, implying that their model displays balanced growth only in the limit where structural change no longer occurs. In contrast, we assume that the capital intensities are the same in goods and services production and are thus able to construct an ABGP that displays structural change.

3. The model of Sposi (2012) also contains structural change in investment. We do not discuss his paper in detail here, because structural change in investment is not the emphasis of his quantitative analysis.

Figure 1
Sectoral composition of GDP, investment, and consumption

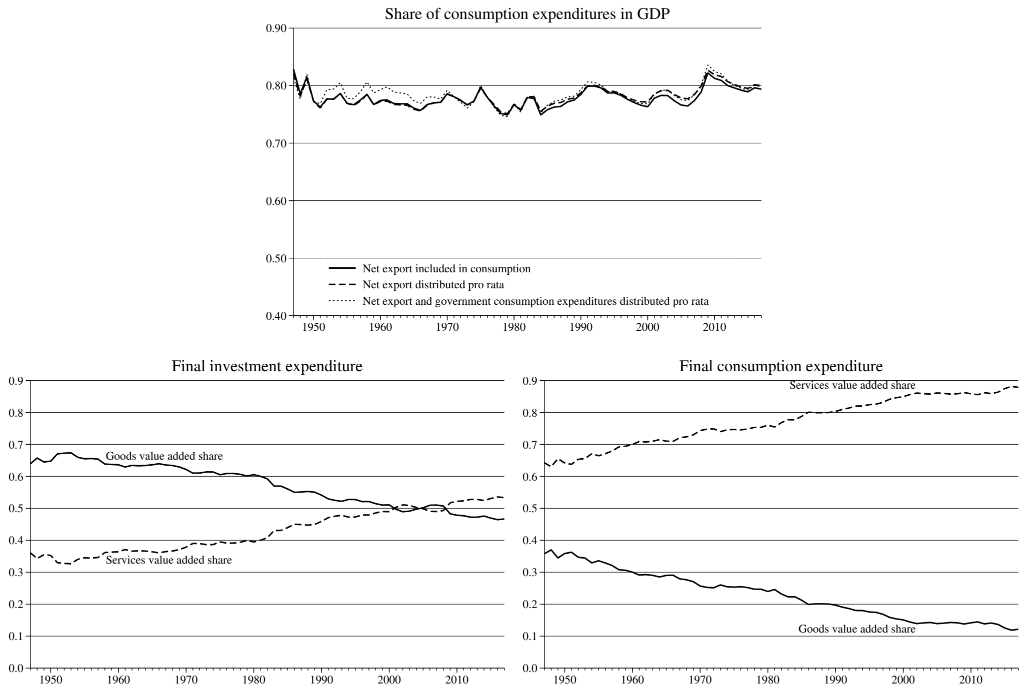

3. EVIDENCE

In this section, we present the structural change facts for the final expenditures on investment and consumption in the US during the post World War II period. This serves to complement existing presentations of the stylized facts of structural change as well as to motivate the framework that we develop in the next section.[4] The basic strategy is to combine US industry data from WORLD KLEMS on industry value added, factor shares and factor inputs with the annual input-output tables from the Bureau of Economic Analysis (henceforth BEA) and then to decompose final expenditure categories into value added components.[5] 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 – investment and consumption – and two value-added components – goods and services.

4. Our evidence complements that presented in Garcia-Santana et al. (2018) 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.

5. WORLD KLEMS covers the period 1947-2014. We extend the WORLD KLEMS data by three years, 2015–2017, using the Integrated Industry-Level Production Account of the BEA and gross output and value added from BEA industry data.

In generating the decompositions, we define the goods sector to consist of agriculture, construction, manufacturing, mining, and utilities.[6] The services sector 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 investment and consumption. We allocate all of net exports to consumption, which has the benefit that the notion of investment and capital in the model will correspond to the notion of domestic investment and capital as measured in the data. Because net exports are not large, the rule for allocating them is not of first-order significance.

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

Figure 1 shows the key facts for this two–by–two decomposition. The upper panel shows that while the share of consumption expenditure in GDP has fluctuated somewhat, it does not have a trend. Our theoretical analysis will focus on aggregate balanced growth paths, and we show that this property holds along such a path. The lower panel shows the shares of goods and services value added in each of final consumption expenditure and final investment expenditure. While investment has a significantly higher goods-valued-added share than does consumption, both investment and consumption 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 in investment is less than the services value added share.[7] 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 in the investment sector and the consumption sector. Third, the value added share of goods differs significantly between investment and consumption, suggesting that it is important that consumption and investment be modelled separately. Because the investment sector uses a disproportionate share of total goods value added, structural change in investment is of particular importance for understanding what happens to the goods producing sector. Its recent decline by a few percentage points had large implications for the level of manufacturing employment, which has received a lot of attention in the policy debate.

Given that relative prices and relative quantities will play an important role in the analysis to come, we also present evidence on each of these. The left graph of the upper

6. 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 was relatively unimportant, we have chosen to combine agriculture and manufacturing into a single goods producing-sector in order to facilitate exposition.

7. As noted in Herrendorf et al. (2013), a different but related problem with the standard assumption that all investment is done in the goods sector is that recently the value of US investment has exceeded the value added produced in the goods sector.

FIGURE 2
Relative prices and relative quantities

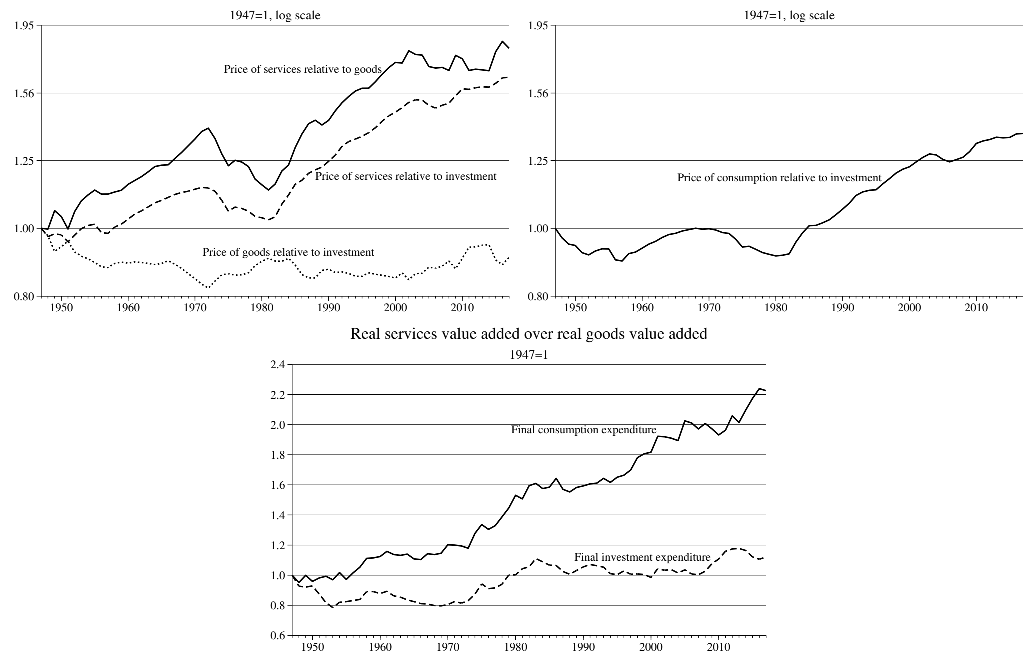

panel of Figure 2 displays time series evidence on three relative prices: the price of services relative to goods, the price of goods relative to investment, and the price of services relative to investment. It reveals that there has been a marked increase in the price of services relative to goods.[8] The behaviour of the price of goods relative to investment has not received much attention in the literature. Interestingly, this relative price has declined overall, with most of the decline occurring in the first 20 years. Subsequent to 1970 it was relatively stable until 2000 and has increased modestly since then. Lastly, the behaviour of the price of services relative to investment is implied by the behaviour of the prices of services relative to goods and of goods relative to investment.

The right graph of the upper panel of Figure 2 displays time series evidence on the price of consumption relative to investment. It reveals that there has also been a marked increase in the price of consumption relative to investment. Notably, the behaviour 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.[9] Given that services are the dominant component of consumption, it is perhaps not too surprising that the price

8. The somewhat unusual behaviour of this relative price in the 1970s is driven by a dramatic spike in the prices of agricultural products and oil. While this suggests that a more detailed analysis might warrant further disaggregation, the somewhat anomalous behaviour of the 1970s should not distract us from the clear secular trend over the entire postwar period.

9. It is important to note that, following the Bureau of Economic Analysis, we include consumer durables as part of consumption. If we reassigned them to investment, then the series for the two prices relative to investment would look somewhat different; their increase would start earlier than in the 1980s and the overall upward trend would be more pronounced; see the figure in Duernecker et al. (2017a).

of consumption relative to investment exhibits behaviour that is qualitatively similar to the price of services relative to investment.

The lower graph of Figure 2 displays time series evidence on the quantity composition of final investment and consumption expenditure, and it reveals two patterns. First, for investment there is no clear trend in the quantity of services relative to goods. Given that the price of services relative to goods has increased considerably, this suggests that a Leontief specification, which leaves the ratio of services to goods constant, will provide a reasonable fit to the composition of final investment expenditure. We will confirm this conjecture in Section 7 below. Second, for consumption there has been a marked increase in the quantity of services relative to goods. Given that the price of services relative to goods has increased considerably, this implies that a CES specification will not be able to replicate this pattern.[10] Instead, one needs income effects that result from non-homotheticities. Below, we will therefore consider a utility specification from the PIGL class which allows for income effects.

4. MODEL

We build a multi-sector extension of the one-sector neoclassical growth model, formulated in continuous time. Motivated by the evidence 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 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 intensity but potentially different TFPs:

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

where $\theta \in (0,1)$ is the common capital-share parameter and A_{jt} is exogenous TFP in sector j. The assumption that the underlying production functions are Cobb-Douglas with identical capital intensities is common in the literature on structural change, as it allows for aggregation and the existence of a balanced growth path. Herrendorf et al. (2015) show that it also does a reasonable job of replicating the sectoral composition in the postwar US.[11]

The outputs of the goods and services sectors are combined to produce final investment using a CES aggregator:

$$X_t = A_{xt} \left(\omega_x^{\frac{1}{\varepsilon_x}} X_{gt}^{\frac{\varepsilon_x - 1}{\varepsilon_x}} + (1-\omega_x)^{\frac{1}{\varepsilon_x}} X_{st}^{\frac{\varepsilon_x - 1}{\varepsilon_x}} \right)^{\frac{\varepsilon_x}{\varepsilon_x - 1}}, \tag{4.1}$$

where $\varepsilon_x \in [0, \infty)$ is the elasticity of substitution between goods and services in the production of final investment and $\omega_x \in [0,1]$ determines the relative weight of inputs from the goods sector in investment production. A_{xt} represents exogenous, investment–specific technical change that is neutral with respect to inputs. Note that the standard case in which investment is entirely produced in the goods sector is captured as the special case in which $\omega_x = 1$.

10. With a CES specification, relative quantities never move in the same direction as relative prices. The most extreme possibility is that relative quantities are constant in the face of relative price changes, as occurs with a Leontief specification.

11. In a different context, Acemoglu and Guerrieri (2008) and Alvarez-Cuadrado et al. (2017) study models in which the two technologies have different capital intensities and show that a balanced growth path exists only asymptotically, when structural change no longer occurs.

One might be concerned that our CES specification is too restrictive because it abstracts from the possibility of non-neutral technical change that differentially affects the inputs from the goods and services sectors. It is important to realize that our CES specification (4.1) encompasses an element of non-neutral technical change, because we do allow for differential rates of technical change in goods and services. As we will see below, this is sufficient to capture the empirical patterns presented in the previous section regarding structural change in the investment sector.

There is an infinitely lived representative household with preferences defined by:

$$\int_0^\infty e^{-\rho t} U(C_{gt}, C_{st}) dt,$$

where $\rho > 0$ is the discount rate. At this point we do not place any restrictions on the form of the instantaneous-utility function $U(\cdot, \cdot)$ aside from the standard ones that it be twice continuously differentiable, increasing, and quasi concave. Below, we will restrict attention to a class of utility functions that includes two familiar examples from the literature, namely, log-CES utility functions and PIGL indirect utility functions.

The household is endowed with one unit of time at each instant and a positive initial capital stock, $K_0 > 0$. Given the form of the utility function, the time endowment is inelastically supplied. 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 requires:

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

5. EQUILIBRIUM PROPERTIES

We study the competitive equilibrium for the above economy.[12] There are representative firms for producing goods, services, and final investment. The household accumulates capital and rents it to the firms. At each point in time there are five markets: three markets for the firm outputs (goods, services, and investment) and two markets for the production factors (labor and capital). The rental prices for capital and labor are denoted by R_t and W_t, the prices for goods and services are denoted by P_{gt} and P_{st}, and the price of final investment is denoted by P_{xt}. As in Ngai and Pissarides (2007), we normalize the price of final investment to one in each period. Arbitrage implies that the interest rate will equal $R_t - \delta$. The exogenous driving forces of our model economy are the three processes for the TFP terms A_{jt}, $j \in \{g, s, x\}$. Given that the equilibrium concept is standard, we do not provide a formal definition of equilibrium.

In the remainder of this section, we provide a partial characterization of the equilibrium. We first derive analytical expressions for the prices of the three outputs in terms of model primitives. As a by product, we also derive expressions for relative expenditure shares of goods and services value added in the investment sector in terms of model primitives. Next we derive an alternative representation of equilibrium production that serves to connect our model to existing versions of the growth model without

12. Since the evidence on relative prices in Figure 2 provides important information regarding parameter values, studying the competitive equilibrium is more natural in our context than studying the planner problem even though the competitive equilibrium will be efficient.

structural change in the investment sector. Lastly, we derive expressions that link equilibrium outcomes with measures of structural change.

5.1. *Output Prices*

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

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

Taking the ratio of the first-order conditions for sector $j \in \{g, s\}$ and rearranging gives:

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

It follows that capital-labor ratios are equalized across sectors. 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 inverses of relative TFPs:

$$\frac{P_{gt}}{P_{st}} = \frac{A_{st}}{A_{gt}}. \tag{5.2}$$

These are all standard results in the structural change literature when the sector production functions are Cobb-Douglas with the same capital intensity; see e.g. Herrendorf et al. (2014).

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

$$P_{xt} = 1 = \frac{\left(\omega_x P_{gt}^{1-\varepsilon_x} + (1-\omega_x) P_{st}^{1-\varepsilon_x}\right)^{\frac{1}{1-\varepsilon_x}}}{A_{xt}}. \tag{5.3}$$

Equations (5.2) and (5.3) are two equations in the two unknowns P_{gt} and P_{st}, and so they allow us to fully characterize the two equilibrium prices in terms of primitives. Straightforward algebra yields:

$$P_{jt} = \frac{A_{xt}\left(\omega_x A_{gt}^{\varepsilon_x-1} + (1-\omega_x) A_{st}^{\varepsilon_x-1}\right)^{\frac{1}{\varepsilon_x-1}}}{A_{jt}}, \quad j \in \{g, s\}. \tag{5.4}$$

The cost minimization problem for the production of final investment also generates a standard expression for relative expenditures on inputs in the production of final investment:

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

This expression describes the nature of structural change in investment as a function of the relative price of goods to services and the elasticity of substitution. Given data on relative prices and relative expenditure shares, this expression can be used to infer values for ω_x and ε_x. We carry out this exercise in Section 7.

Combined with our previous result about the price of goods relative to services in equation (5.2), we can also express relative expenditure shares purely as a function of model primitives:

$$\frac{P_{gt}X_{gt}}{P_{st}X_{st}} = \frac{\omega_x}{1-\omega_x}\left(\frac{A_{st}}{A_{gt}}\right)^{1-\varepsilon_x}. \tag{5.5}$$

It follows that differential, exogenous growth in A_{st} and A_{gt} will lead to structural change in investment as long as ε_x is not equal to one.

5.2. *Alternative Representation of Production*

Our formulation of technology involves a two-level structure in which we start with value-added production functions for goods and services and then produce final investment expenditure by combining valued added from the goods and services sectors. A common alternative formulation in the structural-change literature is to directly posit three value-added production functions, one for each consumption category and another for investment. More concretely, this approach would start with the following production structure:[13]

$$C_{jt} = B_{jt}K_{jt}^{\theta}L_{jt}^{1-\theta}, \quad j \in \{g,s\} \tag{5.6}$$

$$X_t = B_{xt}K_{xt}^{\theta}L_{xt}^{1-\theta}. \tag{5.7}$$

In this formulation, comparing B_{xt} to either B_{gt} or B_{st} provides a measure of the extent to which technical change is investment-biased. Moreover, a simple implication of this more common formulation is that in equilibrium, the price of consumption category $j \in \{g,s\}$ relative to investment is equal to B_{xt}/B_{jt}. It follows that relative TFP growth, and the extent to which technical change is investment-biased, can be inferred directly from data on relative price changes. Additionally, following the steps outlined in Ngai and Pissarides (2007), it is well known how to analyse aggregate balanced growth paths in models with this production structure.

While analytically convenient, the problem with positing this common alternative formulation as a starting point is that it precludes structural change in the investment sector. In what follows we show that although our two-level structure does allow for structural change in investment, in equilibrium, it can be mapped into this more common structure. Identifying the mapping turns out to greatly facilitate the analysis of aggregate balanced growth in the next section. To form some intuition for the mapping, recall that in the last section prices were derived as functions of the primitives of technology. Given prices, one can solve for the cost-minimizing way to produce a unit of investment, i.e., the optimal combination of goods and services. Because goods and services are produced using the same capital-to-labor ratios, this allows one to infer the amount of labor and capital used to produce a unit of investment given the TFPs in the goods and services sectors. Hence, one can infer the relationship between the inputs devoted to the production of investment and the output of the investment sector. The next proposition characterizes this relationship.

Lemma 1. *Along any equilibrium path,* $\forall t \in [0,\infty)$

$$X_t = \mathcal{A}_{xt}K_{xt}^{\theta}L_{xt}^{1-\theta}, \tag{5.8}$$

13. For example, this is the production structure assumed in the review chapter of Herrendorf et al. (2014). We note that the case in which investment is produced in the goods-producing sector is just a special case of this in which $B_{xt} = B_{gt}$.

where

$$\mathcal{A}_{xt} \equiv A_{xt}\left(\omega_x A_{gt}^{\varepsilon_x-1} + (1-\omega_x)A_{st}^{\varepsilon_x-1}\right)^{\frac{1}{\varepsilon_x-1}}, \tag{5.9}$$
$$L_{xt} \equiv \frac{X_{gt}}{A_{gt}K_t^\theta} + \frac{X_{st}}{A_{st}K_t^\theta}, \quad K_{xt} \equiv K_t L_{xt}.$$

Proof. See the Appendix A.

Expression (5.8) has the appearance of a production function, with the term $\mathcal{A}_{xt}$ as TFP. But because it is a relation that holds only in equilibrium we will refer to it as a *pseudo* production function and to $\mathcal{A}_{xt}$ as *effective* TFP. Our model posits Cobb-Douglas production functions at the sector level, whose outputs are aggregated non-linearly to produce final investment. Interestingly, the pseudo production function for investment shows that the Cobb-Douglas property is preserved with regard to the inputs of capital and labor and that the non-linear aggregation only affects the expression for effective TFP.[14]

Note that given the definition of the effective TFP, (5.4) imply that relative prices will be given by:

$$P_{jt} = \frac{\mathcal{A}_{xt}}{A_{jt}}, \quad j \in \{g, s\},$$

which is the analogue of the result that holds in the more common formulation of the production structure in (5.6)–(5.7). Importantly, and a point to which we will return later, prices for the two consumption categories relative to investment are now non-linear functions of all three primitive TFPs in our model.

One can also derive a pseudo production function for aggregate final output, where we define aggregate final output as measured in units of final investment: $Y_t = X_t + P_{gt}C_{gt} + P_{st}C_{st}$.[15] Because payments to inputs will exhaust income for each of the investment- and consumption-producing firms, we have:

$$\begin{aligned} Y_t &= P_{gt}X_{gt} + P_{st}X_{st} + P_{gt}C_{gt} + P_{st}C_{st} \\ &= P_{gt}(X_{gt} + C_{gt}) + P_{st}(C_{st} + X_{st}). \end{aligned}$$

Using that feasibility requires $X_{gt} + C_{gt} = Y_{gt}$ and $X_{st} + C_{st} = Y_{st}$ and 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:

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

14. This result parallels previous results that used CES aggregators and Cobb-Douglas functions with equal shares to derive aggregate relationships; see for example Hsieh and Klenow (2009) and Comin and Hobijn (2010). However, the aggregation properties highlighted in the Lemma have not previously been leveraged to simplify the analysis of dynamics in the context of models of structural change.

15. Duernecker et al. (2017a) stress that this is not how output is defined by the BEA, implying 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. If we wanted to calibrate our model to match the growth of GDP per capita in NIPA, we would need to make sure that the GDP measure is the same in the model and the data.

The previously derived expression for P_{gt} in terms of primitives from equation (5.4) then gives:

$$Y_t = \mathcal{A}_{xt} K_t^{\theta}, \tag{5.10}$$

where $\mathcal{A}_{xt}$ is the effective TFP in the investment sector as defined earlier in (5.9). Note that the aggregate input of labor does not feature explicitly in equation (5.10) because it equals one.

Relation (5.10) is the natural generalization of the relation that one obtains in the more common three-sector model with exogenous sector-specific technical change expressed in equations (5.6)–(5.7). While both our and the more common model have three primitive TFPs, the key difference is that in our framework the level of *effective* TFP in the investment sector is a non-linear function of the three primitive TFP levels. In contrast, in the more common three sector model summarized by equations (5.6)–(5.7), the TFP in the investment sector is itself one of the three primitive TFPs.

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

$$R_t = \theta \mathcal{A}_{xt} K_t^{\theta-1},$$
$$W_t = (1-\theta) \mathcal{A}_{xt} K_t^{\theta}.$$

5.3. *Structural Change*

The dimension of structural change that concerns us is the composition of aggregate economic activity across goods and services. In general, one could measure this by focusing on 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 are given by:

$$\zeta_{jt} \equiv \frac{P_{jt}Y_{jt}}{Y_t} = \frac{P_{jt}Y_{jt}}{P_{gt}Y_{gt} + P_{st}Y_{st}}, \quad j \in \{g, s\}.$$

Substituting the sectoral production functions and using that in equilibrium $K_{gt}/L_{gt} = K_{st}/L_{st} = K_t$ and $P_{gt}A_{gt} = P_{st}A_{st}$, we have:

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

In other words, sectoral nominal value added shares equal sectoral employment shares and we can restrict our attention to characterizing the former.[16]

Since the two value added shares sum to one it is sufficient to study one of them to characterize structural change, and so we focus on the services share ζ_{st}. Carrying out some simple manipulations gives:

$$\zeta_{st} = \underbrace{\frac{X_t}{Y_t}}_{\substack{\text{extensive}\\ \text{margin}}} \cdot \underbrace{\frac{1}{(P_{gt}X_{gt})/(P_{st}X_{st}) + 1}}_{\text{intensive margin}} + \underbrace{\frac{E_t}{Y_t}}_{\substack{\text{extensive}\\ \text{margin}}} \cdot \underbrace{\frac{1}{(P_{gt}C_{gt})/(P_{st}C_{st}) + 1}}_{\text{intensive margin}}, \tag{5.11}$$

16. Labor shares here should be interpreted as shares of efficiency units. If one were to connect the shares of the model to the data, one would need to take into account that shares of raw labor may not be the same as shares of efficiency units.

where we have used that one can write the value added of the services sector and the expenditures on investment and consumption as:

$$\begin{aligned}P_{st}Y_{st} &= P_{st}C_{st} + P_{st}X_{st},\\ X_t &= P_{gt}X_{gt} + P_{st}X_{st},\\ E_t &= P_{gt}C_{gt} + P_{st}C_{st}.\end{aligned}$$

We call the second expression in each term on the right-hand side of (5.11) the *intensive margin* of structural change, as it operates within each final expenditure category. We call the first expression in each term on the right-hand side of (5.11) the *extensive margin* of structural change, as it reflects changes in the composition of final expenditure in terms of investment and consumption. Such changes in the composition of final expenditure will potentially influence the services share for the aggregate economy even if the services expenditure share remains constant within each final expenditure category.

In the next section, we will analyze the existence of an equilibrium that displays the features of balanced growth in the aggregate and structural change at the sectoral level, consistent with the evidence presented in Figure 1. We will show that along such an equilibrium path, the composition of final expenditure between consumption and investment is indeed constant so that only the intensive margin operates and the nature of structural change can be inferred from the behavior of structural change in each final expenditure sector. In contrast, Garcia-Santana et al. (2018) 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.

6. AN ABGP WITH STRUCTURAL CHANGE – THEORETICAL ANALYSIS

As is well known, conventional balanced growth is too strict an equilibrium concept for a multi-sector growth model with structural change, since sectoral labor inputs cannot feasibly grow at constant yet different rates forever. We therefore follow Ngai and Pissarides (2007) and use the less restrictive concept of an *aggregate* balanced growth path, or ABGP for short.

Definition 1. *An ABGP is an equilibrium path along which aggregate variables expressed in units of a numeraire grow at constant, though potentially different rates including zero. In our context with investment as the numeraire, each of K_t, X_t, Y_t, E_t, R_t, and W_t will grow at constant rates (including zero) along an ABGP.*

Importantly, an ABGP allows *sectoral* variables to grow at non-constant rates, thereby permitting structural change to take place "underneath" the ABGP.

6.1. *Existence and Properties of ABGP*

The previous section derived properties that will hold along any equilibrium path for a generic period utility function $U(C_{gt}, C_{st})$. Establishing the existence of an ABGP will require that we place more structure on $U(C_{gt}, C_{st})$. The PIGL class is of particular interest in this context, because, as shown by Boppart (2014), it allows for an ABGP in a model featuring structural change in consumption.[17] Because this class can only

17. Boppart abstracted from structural change in investment

be characterized in terms of its implied indirect utility function, we choose to impose conditions on the indirect rather than the direct utility function.

Given a utility function satisfying the standard regularity conditions, the indirect utility function is defined as:

$$V(P_{gt}, P_{st}, E_t) = \max_{C_{gt}, C_{st}} \left\{ U(C_{gt}, C_{st}) \text{ s.t. } P_{gt}C_{gt} + P_{st}C_{st} \leq E_t \right\}.$$

We will use the following notation:

$$V(t) \equiv V(P_{gt}, P_{st}, E_t), \qquad V_E(t) \equiv \frac{\partial V(P_{gt}, P_{st}, E_t)}{\partial E_t}.$$

As is well known, a "valid" indirect utility function must exhibit several properties. In particular, it must be the case that it is twice continuously differentiable, homogenous of degree zero in prices and total expenditure, non-increasing and quasi convex in both prices. In what follows we will assume that we start with a valid indirect utility function $V(t)$ and consider which additional assumptions are needed for establishing an ABGP with structural change in both investment and consumption.

Our first assumption is about the functional form of the derivative of the indirect utility function with respect to expenditure:

Assumption 1. *$V(t)$ satisfies $V_E(t) = E_t^{\chi-1} P_{st}^{-\chi}$ where $0 \leq \chi \leq 1$.*

While it is perhaps not immediately obvious how to interpret this functional form assumption, we show below that the set of indirect utility functions with this property includes two prominent specifications from the structural change literature: when $\chi = 0$ this condition corresponds to the log-CES specification used by Ngai and Pissarides (2007), and when $\chi > 0$ it corresponds to the PIGL specification used by Boppart (2014).[18]

We continue with additional assumptions that are required to ensure that an ABGP with the desired properties exists. As is well known from the analysis of one-sector models, establishing existence of a balanced growth path will sometimes require an upper bound on the exogenous growth rates of TFP, with the upper bound depending upon preference parameters. That will also be the case here, though we note that there are now multiple sources of exogenous growth and the conditions place joint restrictions on growth rates. Denoting growth rates by "hats", we make:

Assumption 2. *$\forall\, t \in [0, \infty)$:*

$$\chi \left(\frac{\theta}{1-\theta} \widehat{\mathcal{A}}_{xt} + \widehat{A}_{st} \right) < \rho, \tag{6.12}$$

$$0 < \delta + \rho + \chi \left(\widehat{\mathcal{A}}_{xt} - \widehat{A}_{st} \right) + (1-\chi) \frac{\widehat{\mathcal{A}}_{xt}}{1-\theta}. \tag{6.13}$$

Note that this assumption places restrictions on $\widehat{A}_{gt}$ and $\widehat{A}_{xt}$ only through their effect on $\widehat{\mathcal{A}}_{xt}$. Given a value for $\widehat{A}_{st}$, the first equation can be understood as putting a bound on $\widehat{\mathcal{A}}_{xt}$. Recalling that $\chi \geq 0$ and taking as given a value for $\widehat{\mathcal{A}}_{xt}$, the second

18. Given the symmetry between goods and services in the model the reader may wonder why this assumption attributes a distinctive role for services. We postpone discussion of this until the analysis of PIGL utility.

equation can be understood as placing a bound on the growth of $\widehat{A}_{st}$. In the special case of $\chi = 0$, the first condition necessarily holds and the second condition necessarily holds if $\widehat{\mathcal{A}}_{xt}$ is non-negative.

It is standard in the balanced growth literature when looking for a balanced growth path to assume that all exogenous TFP terms grow at constant rates. To this point we have not yet made any such restriction, and we will show later on that imposing it would effectively rule out finding an ABGP that features structural change in investment. But our next assumption does impose a restriction on A_{st}:

Assumption 3. *$\chi\widehat{A}_{st}$ is constant $\forall t \in [0, \infty)$.*

As we show later on, for the log-CES case $\chi = 0$ and the assumption is satisfied independently of the behavior of A_{st}. But for the PIGL case, $\chi > 0$ and the assumption does require that A_{st} grow at a constant rate.[19]

We are now ready to state our main result concerning the existence of an ABGP. Recalling the definition of $\mathcal{A}_x$ in (5.9), we have:

Proposition 1. *Suppose that Assumptions 1–3 hold. Then an ABGP exists if and only if $\widehat{\mathcal{A}}_{xt}$ is constant $\forall t$. Along the ABGP,*

$$\widehat{K} = \widehat{X} = \widehat{Y} = \widehat{E} = \frac{\widehat{\mathcal{A}}_x}{1-\theta}. \tag{6.14}$$

Proof. See Appendix B.

We stress that although the condition of constant growth in $\mathcal{A}_{xt}$ implies a joint restriction on the three primitive TFPs, it should not be interpreted as suggesting that any of the three components A_{xt}, A_{gt}, and A_{st} will adjust endogenously to changes in the other two components so as to ensure the condition holds. Instead, all three components are exogenous variables and the existence condition either does or does not hold. Whether $\mathcal{A}_{xt}$ exhibits constant trend growth is an empirical question that we will answer in the affirmative for the post World War II US in the next section.

6.2. *Structural Change along the ABGP*

Having established conditions for the existence of an ABGP, we now turn to establishing conditions under which the ABGP features aggregate structural change. To this end, it is helpful to restate the decomposition of the share of services value added in aggregate GDP, equation (5.11) from above:

$$\zeta_{st} = \underbrace{\frac{X_t}{Y_t}}_{\substack{\text{extensive}\\ \text{margin}}} \cdot \underbrace{\frac{1}{(P_{gt}X_{gt})/(P_{st}X_{st})+1}}_{\text{intensive margin}} + \underbrace{\left(1-\frac{X_t}{Y_t}\right)}_{\substack{\text{extensive}\\ \text{margin}}} \cdot \underbrace{\frac{1}{(P_{gt}C_{gt})/(P_{st}C_{st})+1}}_{\text{intensive margin}},$$

The extensive margin will not operate along an ABGP, as Proposition 1 established that along an ABGP the investment share, X_t/Y_t, is constant. Turning to the intensive margin, it is useful to restate the equilibrium composition of investment, equation (5.5),

19. The fact that the assumption involves A_{st} rather than A_{gt} is related to the distinctive role of services in Assumption 1.

from above:

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

The next assumption ensures that $(P_{gt}X_{gt})/(P_{st}X_{st})$ decreases in equilibrium, that is, the investment sector of our model displays structural change from goods to services:

Assumption 4. $\varepsilon_x \in [0,1)$*;* $\forall t \in [0,\infty)$ $\widehat{A}_{gt} > \widehat{A}_{st}$*.*

We are left with stating conditions under which we have structural change in the consumption sector. Noting that with our specification we can express time t consumption choices as functions of time t expenditure on consumption and time t prices for goods and services, we assume these functions have the following properties:

Assumption 5. *Goods and services are gross complements:*

$$\frac{\partial C_{it}}{\partial P_{jt}} < 0.$$

Goods are necessities and services are luxuries:

$$\frac{\partial \log(P_{gt}C_{gt})}{\partial \log(E_t)} \leq 1 \leq \frac{\partial \log(P_{st}C_{st})}{\partial \log(E_t)}.$$

The assumption that goods and services are gross complements in consumption implies that the change in relative prices induce structural change also from goods to services in the consumption sector. The assumption that the income elasticity of the expenditure on goods is less than or equal to that of the expenditure on services guarantees that, as expenditure are growing along the ABGP, income effects will not offset the effect of relative prices on the consumption composition. Hence, $(P_{gt}C_{gt})/(P_{st}C_{st})$ decreases.

In sum, under the above conditions, the extensive margin is absent along an ABGP and the intensive margin leads to structural change from goods to services within both sectors. Hence, an ABGP in our model displays aggregate structural change from goods to services:

$$\zeta_{st} = \underbrace{\frac{X_t}{Y_t}}_{\text{constant}} \cdot \underbrace{\frac{1}{(P_{gt}X_{gt})/(P_{st}X_{st})+1}}_{\text{increases}} + \underbrace{\left(1-\frac{X_t}{Y_t}\right)}_{\text{constant}} \cdot \underbrace{\frac{1}{(P_{gt}C_{gt})/(P_{st}C_{st})+1}}_{\text{increases}} \quad \text{increases.}$$

Stated formally, we have shown:

Proposition 2. *Suppose that Assumptions 1–5 hold. Along any ABGP with positive growth in expenditure, there is structural change from goods to services in each sector as well as in the aggregate.*

The results derived so far show that our model can qualitatively account for the stylized facts that we laid out in Figure 1. Specifically, Proposition 1 establishes conditions under which our model has an ABGP along which the share of consumption in GDP is constant, consistent with the evidence presented in the upper panel of Figure 1. Proposition 2 further implies that along an ABGP, the patterns of structural change in investment and consumption are consistent with the secular trends of the nominal value added shares of goods and services in the lower panel of Figure 1.

6.3. *Non-constant Sectoral TFP Growth*

We end this section by deriving the first insight that follows from simultaneously considering structural change in investment and consumption. To this end, we inspect more closely what restrictions the existence condition for an ABGP with structural change imposes on the behavior of sectoral TFP growth.

It is instructive to start with the alternative three-sector production structure (5.6)–(5.7). When looking for an ABGP in this structure, all that is required is that B_{xt} grows at a constant rate that is not too large. Importantly, there is no restriction on the sectoral TFPs B_{gt} and B_{st}; they may grow at the same or different constant rates, but they may also grow at non-constant rates. While our model has the seemingly analogous requirement that constant growth in $\mathcal{A}_{xt}$ is required, this has a very different implication for what is allowed for the primitive sectoral TFPs, as expressed in the next Proposition.

Proposition 3. *Suppose that Assumptions 1–3 hold. A necessary condition for the existence of an ABGP with structural change in investment is that at least one of the growth rates $\widehat{A}_{xt}$, $\widehat{A}_{gt}$, $\widehat{A}_{st}$ is not constant.*

To see why this is the case, we begin with the requirement that

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

must grow at a constant rate. Assume by way of contradiction that all three primitive TFPs grow at constant rates. Given the non-linear term involving A_{gt} and A_{st} on the right-hand side of this expression, the only way that the right-hand side can grow at a constant rate is if one of the following three conditions holds: (i) A_{gt} and A_{st} grow at the same rate, (ii) $\varepsilon_x = 1$, or (iii) ω_x is either 0 or 1. Each of these cases implies that there will be no structural change in investment. In case (i) there is no change in relative prices; in case (ii) there is a unitary elasticity of substitution so that changes in relative prices do not generate any structural change; and in case (iii) investment uses either only goods or only services, thereby ruling out structural change in investment.

While the literature on structural change and balanced growth has tended to focus on constant growth rates of sectoral TFPs, there is no natural reason to favour 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 constant rates does not translate to constant rates of growth in sectoral TFPs. In fact, the previous result shows that not only is it restrictive to consider constant growth rates in sectoral TFPs, but also that imposing this restriction will make it impossible to find an ABGP that exhibits structural change.

In a later section, we examine the behavior of $\mathcal{A}_{xt}$ and its various components in the data. Interestingly, we will find that it grows at an approximately constant average rate despite the fact that this is not true for each of the components. This highlights the importance of understanding the conditions under which aggregate balanced growth can emerge despite the fact that growth rates in sectoral TFP do not grow at constant rates on average.

6.4. *Utility Functions that Permit an ABGP with Structural Change*

In this subsection, we study two classes of utility functions that are covered by our previous results: the log-CES case considered by Ngai and Pissarides (2007) and the

PIGL case considered by Boppart (2014).

6.4.1. Log-CES Utility. We start by establishing that the log of homothetic CES utility satisfies Assumption 1.[20] We also establish that this is not the case for the non-homothetic CES that was applied in the context of structural change by Comin et al. (2018). We assume that $U(C_{gt}, C_{st}) = \log C_t$ where C_t is implicitly defined by:

$$C_t = \left(\omega_c^{\frac{1}{\varepsilon_c}} C_{gt}^{\frac{\varepsilon_c - 1}{\varepsilon_c}} + (1-\omega_c)^{\frac{1}{\varepsilon_c}} C_t^{\frac{\sigma-(\varepsilon_c-1)}{\varepsilon_c}} C_{st}^{\frac{\varepsilon_c - 1}{\varepsilon_c}} \right)^{\frac{\varepsilon_c}{\varepsilon_c - 1}}. \tag{6.15}$$

$\omega_c \in [0,1)$ is a relative weight, $\varepsilon_c \in [0,1)$ is the elasticity of substitution, and $\sigma > 1$ is the parameter determining whether or not the CES is homothetic. All three parameters are constant.

In the special case of $\sigma = \varepsilon_c - 1$, C_t disappears from the right-hand side and we have the standard homothetic CES:

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

In contrast, in the general case of $\sigma \neq \varepsilon_c - 1$, there is no closed-form solution for C_t and (6.15) *implicitly* defines a non-homothetic CES utility. If $\sigma > \varepsilon_c - 1$, then the relative weight on services increases as C_t increases, implying that the income elasticity of services is larger than one. If $\sigma < \varepsilon_c - 1$, then the relative weight on services decreases as C_t decreases, implying that the income elasticity of services is smaller than one.

Solving for the expenditure shares of goods and services as in Comin et al. (2018) and substituting the results back into the utility function (6.15), we obtain an implicit expression for the indirect utility function of the non-homothetic CES class:

$$V(P_{gt}, P_{st}, E_t) = \log(E_t) - \frac{1}{1-\varepsilon_c} \log \left(\omega_c P_{gt}^{1-\varepsilon_c} + (1-\omega_c) V(P_{gt}, P_{st}, E_t)^{\sigma-(\varepsilon_c-1)} P_{st}^{1-\varepsilon_c} \right).$$

Using standard arguments one can show that the first term on the right-hand side implies that indirect utility is increasing in E_t holding prices constant. And the second term on the right-hand implies that indirect utility is decreasing in both prices holding E_t constant. The second term can be understood as the price index that the individual faces. Importantly, the non-homotheticities imply that the level of expenditure influences the price index in a somewhat complicated fashion. In particular, it is apparent that, in general, $V_E(t) \neq E_t^{\chi-1} P_{st}^{-\chi}$ and Assumption 1 will not be satisfied. This is not too surprising given the fact that Comin et al. (2018) are only able to derive an asymptotic balanced growth path.[21]

However, in the special case when $\sigma = \varepsilon_c - 1$, $V_E(t) = E_t^{-1}$ and Assumption 1 is satisfied with $\chi = 0$. Hence we can apply Proposition 1 in this case to conclude that an ABGP exists if and only if $\widehat{\mathcal{A}}_{xt}$ is constant, and that the common growth rate along the ABGP is given by $\widehat{\mathcal{A}}_x/(1-\theta)$. Note that for the case when $\chi = 0$, the first condition in Assumption 2 holds automatically and the second condition is satisfied if $\widehat{\mathcal{A}}_x$ is positive,

20. Having $\log C_t$ instead of C_t is crucial for establishing the results to follow. If one instead assumed an intertemporal elasticity of substitution that is constant but different from one, then an ABGP would not exist.

21. Note that for similar reasons, Assumption 1 does not hold for the generalized Stone-Geary specification that Buera and Kaboski (2009) and Herrendorf et al. (2013) used. Again, the non-homotheticities imply that the level of expenditure influences the price index in a somewhat complicated fashion and only an asymptotic balanced growth path exists.

which is equivalent to saying that growth measured in units of the investment good is positive along the ABGP. As in the standard one-sector growth model, one does not need to put an upper bound on TFP growth when period utility is given by $\log C_t$. Lastly, when $\chi = 0$ Assumption 3 is also trivially satisfied.

It is instructive to briefly compare our model with log-CES utility to the model with the alternative production structure (5.6)–(5.7) and log-CES utility. The latter is essentially the model of Ngai and Pissarides (2007) with two consumption goods. This alternative model has an ABGP with structural change if and only if $\widehat{B}_{xt}$ is constant, with the common growth rate along the ABGP given by: $\widehat{B}_x/(1-\theta)$. It is remarkable how similar the results are, despite the fundamental differences between the two models. The reason for this is that, as shown by Lemma 1, the investment sector of our model has a pseudo production function with the same functional form as the production function of the alternative model. The one crucial difference is that in our model the effective TFP $\mathcal{A}_{xt}$ is a function of the primitives A_{xt}, A_{gt}, and A_{st}, whereas in the alternative model B_{xt} is itself a primitive.

Turning now to the applicability of Proposition 2, we note that Assumption 4 is independent of the utility function and Assumption 5 is satisfied as long as $\varepsilon_c \in [0, 1)$. Because the log-CES utility function is homothetic, the income elasticities of goods and services are both equal to one. Therefore, Proposition 2 holds for log-CES utility and the ABGP displays structural change from goods to services. Moreover, following the same steps as above for the derivation of (5.5) for the expenditure shares in investment, we obtain an analytical solution for the ratio of expenditure shares in consumption, which is exactly as in Ngai and Pissarides (2007):

$$\frac{P_{gt}C_{gt}}{P_{st}C_{st}} = \frac{\omega_c}{1-\omega_c}\left(\frac{A_{st}}{A_{gt}}\right)^{1-\varepsilon_c}. \tag{6.16}$$

6.4.2. PIGL Indirect Utility. We now consider the class of PIGL utility functions, which captures non-homotheticities while permitting for an ABGP with structural change. Note that the PIGL class does not encompass the previous homothetic class of log-CES utility functions as a special case. An exception to that statement is the degenerate Cobb-Douglas case, which is a member of both classes of utility functions.

Since the PIGL class does not have a known direct utility representation in general, we will work with indirect utility. We start from the same functional form and parameter restrictions as Boppart (2014):

$$V(E_t, P_{gt}, P_{st}) = \frac{1}{\chi}\left(\frac{E_t}{P_{st}}\right)^{\chi} - \frac{\nu}{\gamma}\left(\frac{P_{gt}}{P_{st}}\right)^{\gamma} - \frac{1}{\chi} + \frac{\nu}{\gamma}, \tag{6.17}$$

with $\nu > 0$ and $1 > \gamma > \chi > 0$. Note that the elasticity of substitution between goods and services does not equal γ and is not even constant; see Boppart (2014) for the exact formula. And the implied intertemporal elasticity of substitution equals χ, and therefore is not restricted to equal one as was the case with log-CES utility. The assumption that $\chi > 0$ implies that services have an income elasticity greater than unity, i.e., they are a luxury. To see this, use Roy's Identity to derive the expenditure share of goods:

$$\frac{P_{gt}C_{gt}}{E_t} = \nu\left(\frac{E_t}{P_{st}}\right)^{-\chi}\left(\frac{P_{gt}}{P_{st}}\right)^{\gamma}. \tag{6.18}$$

It is apparent that with $\chi > 0$, the expenditure share for goods is decreasing in E_t, which implies that the expenditure share for services is increasing in E_t.

As shown by Boppart (2014), given our parameter restrictions the functional form in (6.17) represents a valid indirect utility function as long as the parameters satisfy:

Assumption 6.

$$E_t^{\chi} \geq \frac{1-\chi}{1-\gamma}\nu P_{gt}^{\gamma} P_{st}^{\chi-\gamma} \quad \forall t \in [0,\infty). \tag{6.19}$$

Condition (6.19) is somewhat unappealing since it is formulated in terms of endogenous variables. Later on we estimate ν, χ, and γ and show that it is satisfied in every period using postwar US data for E_t, P_{gt}, and P_{st}. But in this section we are interested in establishing the existence of an ABGP, and so we will need that the condition holds at every point along an ABGP. It will be more transparent to supply sufficient conditions after we have characterized some properties that any ABGP must have for these preferences. For now, we therefore assume that the condition holds and later on we present sufficient conditions for that to be the case.

To apply Proposition 1 in the PIGL case, we must ensure that Assumption 1 holds. Differentiating (6.17) with respect to E_t gives that $V_E(t) = E_t^{1-\chi} P_{st}^{\chi}$. Since we assumed that $1 > \gamma > \chi > 0$, Assumption 1 holds. Proposition 1 therefore applies to the PIGL case and an ABGP exists if and only if $\widehat{\mathcal{A}}_{xt}$ is constant. The common growth rate along the ABGP is again equal to $\widehat{\mathcal{A}}_x/(1-\theta)$. To apply Proposition 2 in the PIGL case, recall that Assumption 4 is independent of the utility function. Boppart (2014) showed that Assumption 5 holds given our assumptions on parameter values. Hence, Proposition 2 implies there is structural change from goods to services.

Importantly, since $\chi > 0$, Assumptions 2 and 3 are now somewhat more stringent than in the log-CES case. In particular, Assumption 3 now imposes that $\widehat{A}_{st}$ is constant. And, as noted before, the first condition in Assumption 2 imposes a joint upper bound on the two constant growth rates, $\widehat{A}_s$ and $\widehat{\mathcal{A}}_x$. The second condition can be interpreted as imposing an upper bound on $\widehat{A}_s$ given a value of $\widehat{\mathcal{A}}_x$.

Having established that $\widehat{A}_{st}$ and $\widehat{\mathcal{A}}_{xt}$ must both be constant along an ABGP in this case, we can now provide a sufficient condition for Assumption 6 to hold along an ABGP. It is useful to rewrite condition (6.19) as:

$$\frac{1-\gamma}{1-\chi} \geq \nu \left(\frac{E_t}{P_{st}}\right)^{-\chi} \left(\frac{P_{gt}}{P_{st}}\right)^{\gamma} \quad \forall t \in [0,\infty). \tag{6.20}$$

The left-hand side of (6.20) is a constant and lies between zero and one, given that we assumed $1 > \gamma > \chi > 0$. The right-hand side of (6.20) is just the expenditure share for goods displayed above. Starting with the second term on the right-hand side, it follows from our earlier analysis that

$$\widehat{E} - \widehat{P}_s = \frac{\widehat{\mathcal{A}}_x}{1-\theta} - (\widehat{\mathcal{A}}_x - \widehat{A}_s) = \frac{\theta \widehat{\mathcal{A}}_x}{1-\theta} + \widehat{A}_s.$$

A sufficient condition for this to be positive is that both growth rates are positive. We then have that $(E_t/P_{st})^{-\chi}$ will decrease monotonically to zero along an ABGP. Turning to the third term on the right-hand side, equation (5.2) showed that $P_{gt}/P_{st} = A_s/A_{gt}$. Since Assumption 4 implies that A_s/A_{gt} is strictly decreasing, the third term is strictly decreasing. Combining the two results, a sufficient condition for the right-hand side of equation (6.20) to be decreasing is that $\widehat{A}_{st}$ and $\widehat{\mathcal{A}}_{xt}$ grow are positive and constant and $\widehat{A}_{gt} > \widehat{A}_s$ for all t. Under these conditions, Assumption 6 will hold at each point along an ABGP as long as it holds in the initial period. Ensuring that it holds in the initial period basically amounts to a choice of units.[22]

In Subsection 7.2 below, we demonstrate that the non-homotheticicities permitted by the PIGL utility are crucial for matching the shares of goods and services in US consumption expenditure during the postwar period. While this result is very related to the results in Boppart (2014), we note that our formulation uses value added shares whereas his formulation used final expenditure shares in consumption. More generally, he did not consider structural change in investment.

We close the current subsection by addressing a feature of the indirect utility function in (6.17) that is somewhat curious, namely, expenditure are expressed in units of services, which is related to the earlier remark concerning the apparent distinctive role of services in Assumption 1. One might wonder about an alternative specification that expresses expenditure in units of goods while choosing $\chi < 0$ in order to maintain that services are a luxury. If we were to adopt this alternative, it would lead us to reverse the roles of goods and services in Assumptions 1–3, which would give rise to an alternative set in parameter space that permits an ABGP. The reason that we have not formally pursued this alternative is that it is fragile in that for empirically reasonable parameter values the indirect utility function is not valid and the implied expenditure shares are not well defined. To see why, note that if we repeated our analysis for the alternative specification, instead of (6.18), we would obtain:

$$\frac{P_{st}C_{st}}{E_t} = \nu \left(\frac{E_t}{P_{gt}}\right)^{-\chi} \left(\frac{P_{st}}{P_{gt}}\right)^{\gamma}. \tag{6.21}$$

An ABGP would now require that E_t/P_{gt} grow at a constant rate, and if there is positive growth along the AGBP, then E_t/P_{gt} would have to grow at a positive constant rate. Assume in addition that P_{st}/P_{gt} does not decline, which is certainly in line with the existing empirical evidence. In this case, given χ would now be negative, the right-hand side of the above expression would be bounded from below by an expression that grows at a constant positive rate:

$$\nu \left(\frac{E_t}{P_{gt}}\right)^{-\chi} \left(\frac{P_{st}}{P_{gt}}\right)^{\gamma} \geq \nu \left(\frac{E_t}{P_{gt}}\right)^{-\chi} \left(\frac{P_{s0}}{P_{g0}}\right)^{\gamma}.$$

Hence, the left-hand side of (6.21) would now grow without bound. This would imply that at some point the expenditure share would exceed one and the equivalent form of Assumption 6 would be violated. It follows that the alternative specification does not permit an analysis of the asymptotic behavior along an ABGP for empirically reasonable parameter values. Since in Section 8 below, we will establish that our model has important implications also asymptotically, we want a specification that is well defined over the whole time horizon $[0, \infty)$. Therefore, we do not pursue the alternative specification further.[23]

22. To see this note that because the right-hand side of equation (6.20) converges to zero, the condition must be satisfied at some t. Therefore, we can just re-normalize the TFPs so that these values are the period 0 values.

23. Note that the same issue applies to the analysis of Boppart (2014), and like us, he focused on the specification in which expenditure is expressed relative to the price of services. Note too that the issue concerns only the asymptotic behavior along an ABGP. Given that expenditure shares in the data are naturally in $[0, 1]$, there would not be a problem with using the alternative specification to fit the observed expenditure shares in postwar US data.

FIGURE 3
Shares in investment

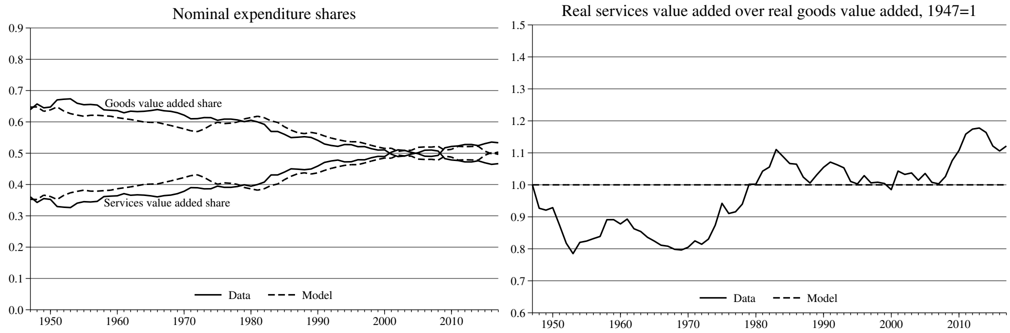

7. ABGP AND STRUCTURAL CHANGE – EMPIRICAL ANALYSIS

In this section, we complement the preceding theoretical analysis with an empirical assessment of two questions. First, to what extent does our framework capture the secular trends in the goods and services expenditure shares in both the investment and consumption expenditure data for the US that we displayed in Figure 1? Second, to what extent do we see roughly constant growth in $\mathcal{A}_{xt}$ in the US data as required for aggregate balanced growth?

7.1. *Empirical Properties of the Investment Aggregator*

Our goal in this subsection is to assess whether the assumed log-CES structure for the investment aggregator is empirically reasonable. To do this, we ask whether there are values for ω_x and ε_x such that, when taking relative prices as given by the data, the log-CES structure is able to capture the key secular changes in relative expenditure shares.[24] In particular, we start with equation (5.5) for the relative expenditure shares within the final investment expenditure sector, which we repeat here for convenience:

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

Using the data as presented in Section 3, we calibrate ω_x and ε_x to minimize the sum of squared deviations between the implied relative expenditure shares and the relative 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.

Figure 3 shows the fitted values for the expenditure shares and the ratio of the real quantities along with their corresponding values in the data. We find that the fitted values track the secular change in the expenditure shares of the investment sector very well. In particular, the model captures 86% of the increase in the data in the expenditure share for services between the first and the last period. The implied parameter values are $w_x = 0.65$ and $\varepsilon_x = 0.00$. The value of zero for the elasticity parameter implies

24. We note that this exercise does not impose that the data lie along an ABGP.

that the aggregator is Leontief. While it may seem surprising that the aggregator in the investment sector features such a low degree of substitutability, this result was actually foreshadowed by the evidence presented in Section 3. In particular, we showed that the real ratio of services to goods in investment had no trend despite large changes in relative prices. With CES aggregators and complements, real input ratios move in the opposite direction of expenditure shares. That is, with an elasticity of substitution less than unity but larger than zero, 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 caused by the change in prices. In the limiting case of a Leontief aggregator, there is a constant ratio of real quantities. Therefore, a Leontief specification accounts best for the behavior in the data, leaving the quantity of services relative to goods constant.

7.2. *Empirical Properties of the Consumption Aggregators*

In this subsection we repeat the analysis of the previous section, but now considering the model's ability to account for structural change within the consumption sector. In particular, we return to the two special cases of interest that are covered by our previous results: the log-CES case and the PIGL case.

We begin with the case of log-CES utility and repeat the same step that we took to estimate the parameters of the investment aggregator. In particular, we take prices as given from the data and ask what values of ω_c and ε_c provide the best fit to the relative expenditure shares. The result of this exercise is $\omega_c = 0.29$ and $\varepsilon_c = 0.00$. Note that the weight on goods in the consumption aggregator is considerably less than the weight on goods in the investment aggregator obtained in the previous subsection. The result that a Leontief specification provides the best fit is perhaps not surprising given the previous work of Herrendorf et al. (2013), as they estimated a Leontief specification based on a three-sector specification in which the goods sector was split into agriculture and non-agriculture.

The upper panel of Figure 4 shows the fit of the estimated model to the data. While the estimated model generates a significant change in the expenditure share for services between the first and the last period, it accounts for only about half of the observed change. This result is to be expected: in Section 3, we documented that the ratio of real services to real goods within the consumption sector increases, and as previously noted a CES specification cannot generate this outcome from an increase in the relative price of services. This result is also consistent with recent work by Boppart (2014) and Comin et al. (2018), who both argue that non-homotheticities are essential for matching structural change in consumption.

Next we repeat this exercise for the case of PIGL utility, which does allow for non-homotheticities. For this specification, there are three parameters to estimate: χ, γ, and ν. The parameter profile that minimizes the sum of squared deviations of the expenditure shares predicted by the model from the data is $\nu = 0.44$, $\chi = 0.55$, and $\gamma = 0.69$. As noted previously, given these parameter values, Assumption 6 holds period by period on the postwar US data. Hence, the indirect utility function is well defined.

The lower panel of Figure 4 shows the fit of the estimated PIGL specification to the data. The PIGL specification does an excellent job of capturing the secular trends in the data, and it does much better than the log-CES specification. In particular, it captures 86% of the change in the expenditure shares between the first and the last period. Given

Figure 4
Shares of consumption

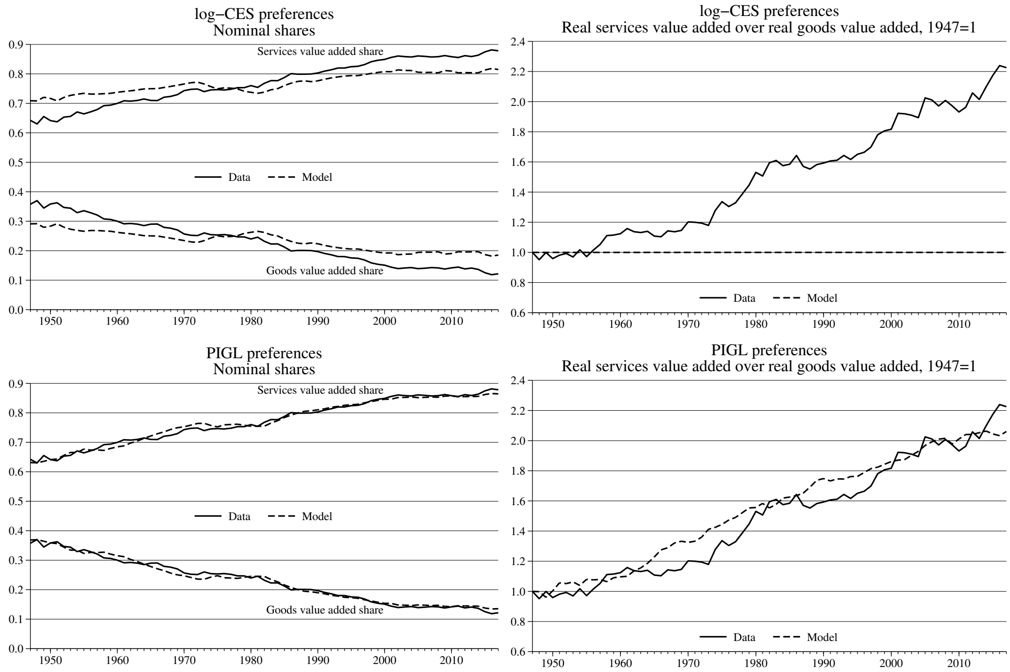

the work of Boppart (2014), this is perhaps not surprising, but, as previously noted, the two exercises are not identical; our estimation is based on matching value added shares in NIPA data whereas his was based on matching final consumption expenditure shares in micro data.[25]

It is of interest to assess the relative contribution of income versus relative price effects from the calibrated PIGL specification. Figure 5 suggests that under the PIGL specification a larger part of structural change (roughly sixty percent) is due to the income effect than is due to the price effect (roughly forty percent), although after 1980 both effects are similarly important. To establish this, the figure plots the implied changes in expenditure shares when one of the channels is present and the other one is shut down. To assess the contribution of income changes to structural change, we use the calibrated PIGL specification, hold relative prices at their 1978 values, and calculate the expenditure shares implied by the observed changes in per capita consumption expenditure. To assess the contribution of relative price changes to structural change, we hold per capita consumption expenditure at their 1978 values and calculate the

25. Boppart finds a somewhat smaller role for the income effect than we do (his estimate of $\chi = 0.22$ compared to our estimate of $\chi = 0.56$). This may seem surprising in light of the result in Herrendorf et al. (2013) that final-expenditure data, which Boppart uses, imply larger income effects than value-added data, which we use. However, Herrendorf et al. (2013) work with aggregate NIPA data, as we do here, whereas Boppart uses micro data. This prevents a simple comparison of estimates. For example, whereas NIPA captures the dramatic increase in expenditure associated with government provided health and education services, they do not show up in household expenditure surveys.

Figure 5

Contribution of changes in income and relative prices under PIGL utility

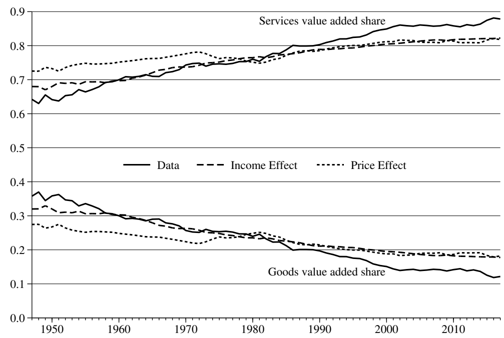

expenditure shares implied by the changes in the relative prices. We choose 1978 as our reference year to facilitate the comparison of the income and price effects, and also the comparison with the log–CES preferences where the predicted expenditure shares cross the actual expenditure shares only once in 1978.

We conclude this subsection with some comments regarding the choice between the log-CES and PIGL specifications. We have seen that the PIGL indirect utility function allows us to best match the observed time series for structural change within consumption. This suggests that one should use the PIGL specification when the empirical fit of the model is the main focus. Nonetheless, for at least two reasons the log-CES specification remains an important benchmark specification whose properties one needs to understand. First, as we show in the next section, only with CES utility does our multi-sector model reduce to the two-sector model with investment and consumption. It is important to establish this link, because it shows what restrictions one needs to impose when using the very popular two-sector model. Second, as we show in the next section, only with CES utility do we have an analytical expression for the model-implied price of consumption relative to investment, P_{ct}. This is useful when one is interested in studying structural change in the presence of investment-biased technical change as reflected in P_{ct}.

7.3. *Empirical Analysis of Technical Progress*

In this subsection, we use standard growth accounting methods as in Solow (1957) to estimate the three primitive TFP terms A_{gt}, A_{st}, and A_{xt}. We then construct an empirical measure of $\mathcal{A}_{xt}$ and assess the theoretical conditions required to guarantee the existence of an ABGP.

We proceed in several steps. In the first one, we solve for A_{jt}, we set $A_{j0} = 1$ and calculate the growth rates by using the information provided by WORLD KLEMS on real value added, capital and labor services, and factor shares:

$$\widehat{Y}_{jt} = \widehat{A}_{jt} + \theta \widehat{K}_{jt} + (1-\theta)\widehat{L}_{jt}, \quad j \in \{g, s\},$$

where θ is the aggregate capital share averaged over the entire period. Since all terms except for $\widehat{A}_{jt}$ are observable, we can solve for $\widehat{A}_{jt}$ and, given the normalization $A_{j0} = 1$,

FIGURE 6
Sectoral TFPs

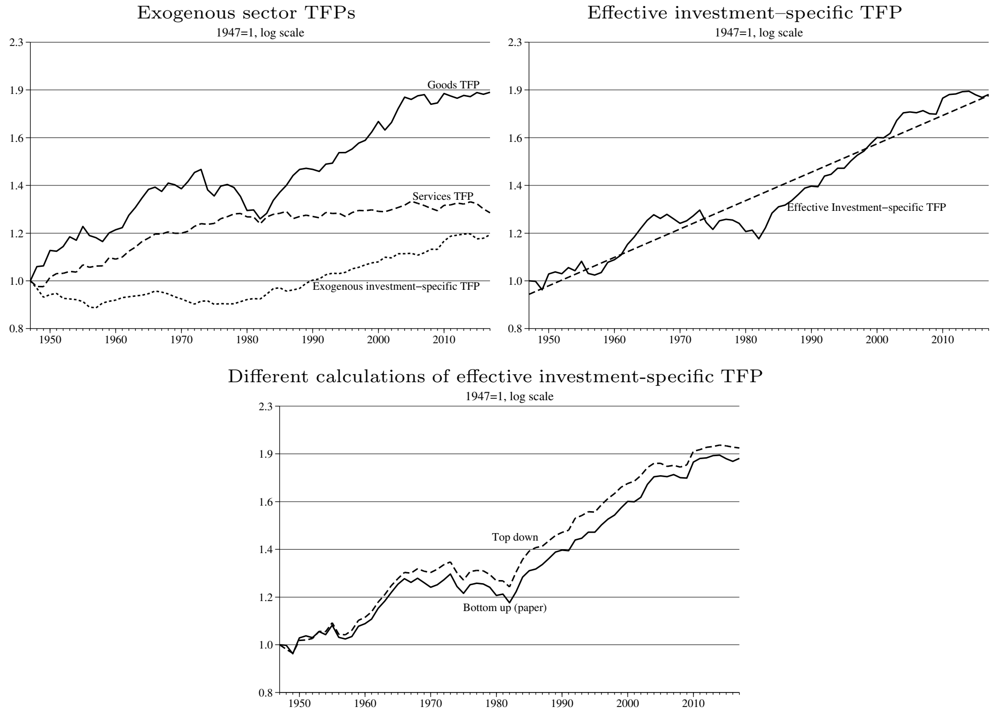

construct A_{jt}.[26]

In the second step we compute A_{xt}. In particular, we again set $A_{x0} = 1$ and use:

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

From this, we compute the implied growth rate series for $\widehat{A}_{xt}$ given that all other components of the identity are observable.[27] Note that in this step we did not impose our functional form for the investment aggregator.

The upper-left graph of Figure 6 shows the implied series for the level of each of the three primitive TFP terms relative to their initial values. Several features are worth noting. First, the trend behavior of A_{xt} seems to have changed over time; it is relatively flat between 1947 and 1980 but exhibits modest and relatively constant positive growth since 1980. The series for A_{gt} seems to exhibit medium-run fluctuations around a fairly constant trend line. Of particular interest is the decade of the 1970s, which we commented on earlier in Section 2. Shocks to oil and food prices specific to this period led to large

26. Alternatively, we could have allowed for each sector to have a different but time invariant capital share given by the sectoral time average, or to have time and sector variation. Given that our focus is on secular changes, these alternatives give very similar answers and so we use the simplest specification as our benchmark.

27. In this calculation, we approximate the continuous–time Divisia index with the Törnqvist index from the data.

temporary deviations. The series for A_{st} also exhibits changing growth rates over time. Between 1947 and the mid-1960s, this series exhibits modest positive growth at roughly 0.8% per year, but in the period after the mid-1960s the growth rate drops to about 0.2% per year. We will return to this behavior later in this section. Lastly, we note that even in the post-1980 period in which growth in A_{xt} increases, the growth in A_{gt} is more than three times larger than the growth in A_{xt}.

Having calculated all three of the primitive TFP series A_{xt}, A_{gt}, and A_{st} from the data, and using our earlier estimates for ω_x and ε_x, we can solve for the implied series for $\mathcal{A}_{xt}$:

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

The upper right panel of Figure 6 shows the time series behavior of $\mathcal{A}_{xt}$. Given our previous analysis, we are very interested in the trend behavior of this series and so have also included the best fitting constant growth line. Although there are some significant departures from the trend line, most notably in the 1970s when temporary shocks had a large impact on TFP in the goods sector, the constant growth line does a reasonable job of tracking the trend behavior in this series. That is, despite the changing behavior of the growth rates of some of the primitive TFPs over time, this series exhibits approximately constant growth, consistent with one of the conditions required for existence of an ABGP.

We emphasize that our construction of the $\mathcal{A}_{xt}$ series did not directly use information about approximately constant average growth of GDP in the US data. That is, we used a "bottom-up" approach to computing $\mathcal{A}_{xt}$, building it up from its underlying components. Alternatively, we could have used a "top-down" approach by applying growth accounting methods directly to our pseudo-aggregate production function, noting that, in the context of our model, we measure output in units of the investment good. The results of this exercise are shown in the bottom panel of Figure 6. It is reassuring that the two approaches yield very similar series for $\mathcal{A}_{xt}$.

One may wonder how our result of approximately constant growth in $\mathcal{A}_{xt}$ and an ABGP can be reconciled with studies that conclude the US is experiencing a growth slowdown. A key point in this context is that our ABGP features constant growth of output measured in units of investment, whereas official measures of output growth rely on a Fisher index. Figure 7 displays these two distinct measures from the data. Interestingly, the series for output measured in units of investment does not display the growth slowdown that is evident in the NIPA measure; see Duernecker et al. (2017a) for a more detailed analysis of this issue.

In our theoretical analysis we imposed several assumptions that involved technology and preference parameters. We can now revisit these assumptions in light of our estimates. When χ is positive, as in our estimated PIGL specification, Assumption 3 requires that A_{st} grows at a constant rate. As discussed earlier in the context of Figure 6, A_{st} has displayed approximately constant growth since around 1965, but the growth rate in the pre-1965 period was somewhat larger. To be specific, the average growth rate of A_{st} is 0.8% in the pre-1965 period and 0.2% in the post-1965 period. Given that the growth rate of $\mathcal{A}_{xt}$ is approximately constant over the entire period, our analysis suggests that one should have two distinct ABGPs, one for the pre-1965 period and one for the post-1965 period. Our analysis allows us to assess the quantitative magnitude of this difference. Importantly, holding the growth rate of $\mathcal{A}_{xt}$ constant, a change in the growth rate of A_{st} does not affect the growth rate along the ABGP (see Proposition 1). Rather, a change in the growth rate of A_{st} simply affects the level of output by affecting the rate of return to capital along the ABGP (see Equation (B22) in the Appendix). Given our estimate

FIGURE 7

GDP per capita growth

Numeraire investment versus Fisher index

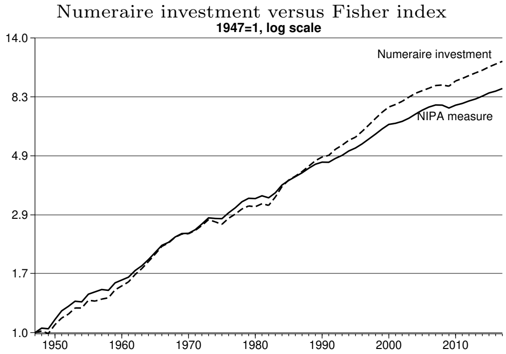

of $\chi = 0.55$, a decrease in the growth rate of A_{st} by .006 increases the rental rate of capital by .0034. Assuming that $\theta = 1/3$, the implied level effect on total output is approximately 1%.[28] We conclude that the change in the growth rate of A_{st} between the pre- and post-1965 periods is quantitatively not of first-order importance.

Our estimation has focused on the static preference parameters and the sectoral TFP series; to evaluate the conditions in Assumption 2 for the case of PIGL, we will assume standard values for the remaining parameters: $\theta = 1/3$, $\delta = 0.08$, and $\rho = 0.04$. Both conditions of Assumption 2 are satisfied in every period when we use these values together with our estimated preference parameters and our estimated average TFP growth rates.

We previously remarked that Assumption 6 is satisfied every period using our estimated values for ν, γ and χ and the values from the data for expenditure and prices. Lastly, although $\widehat{A}_{gt} > \widehat{A}_{st}$ does not hold period by period, it does hold when using average growth rates over the period. From the perspective of understanding the secular patterns of structural change, it is the average growth rates over longer periods that are of interest.

One insight from our analysis that we previously highlighted was that one cannot find an ABGP if one imposes that all primitive TFPs grow at constant rates. In the next two sections, we highlight two additional insights from our analysis.

8. DO SERVICES TAKE OVER THE ECONOMY ASYMPTOTICALLY?

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. This is often referred to as Baumol's cost disease. Taken to the extreme, his argument suggests that if there is one sector with low (possibly zero) productivity growth, it will eventually dominate the entire economy.

Importantly, Baumol's formal analysis abstracted from capital. Ngai and Pissarides (2007) built a model of structural change with capital and showed that it has a different asymptotic implication. In their model, only one sector is assumed to produce capital.

28. Assuming an annual growth rate of 2%, depreciation of 8% and discounting of 4%, the value of R would increase from 0.14 to 0.1434. This implies that the level of output in the new ABGP relative to the original is equal to $(.1434/.1400)^{\theta/(\theta-1)} = 0.988$.

Since the capital-to-output ratio is constant along an ABGP, the capital-producing sector does not disappear, even asymptotically. If one identifies the capital-producing sector with manufacturing, which has higher productivity growth than services, then their model implies that the sector with the slowest productivity growth will dominate consumption production, but not the entire economy.

To see how this result is related to our work, it is useful to go back to our alternative production structure (5.6)–(5.7) with one investment sector and consumption sectors for goods and services. A version of the result of Ngai and Pissarides (2007) obtains in this model if we assume that the investment sector and the goods sector have higher TFP growth than the services sector and that goods and services are gross complements. While the service sector will asymptotically take over the consumption sector, the investment sector produces a constant fraction of GDP along an ABGP. Hence, the share of services will converge to that of consumption, and services will not take over the economy in the limit.

Now consider instead our more general specification in which both goods and services are used to produce investment. Additionally, assume again that TFP grows more strongly in the goods sector than in the services sector (including in the limit), and that goods and services are gross complements in producing both consumption and investment. Along an ABGP, it remains true that investment will be a constant fraction of output. But importantly, the services expenditure share in the investment sector will now converge to one too. The next proposition summarizes this result:

Proposition 4. *If $\varepsilon_x \in [0,1)$, $\widehat{A}_{gt} > \widehat{A}_{st}$, $\lim_{t\to\infty} \widehat{A}_{gt} > \lim_{t\to\infty} \widehat{A}_{st}$, then along any ABGP:*

$$\lim_{t\to\infty} \frac{P_{st}X_{st}}{X_t} = 1.$$

We note that, given our earlier assumptions, for both log-CES and PIGL utility, our model also generates the standard result that: $\lim_{t\to\infty} P_{st}C_{st}/E_t = 1$. Therefore, it follows that services will take over the entire economy asymptotically in both cases.[29]

Because these results hold along an ABGP in which output measured in units of investment grows at a constant rate, one might be tempted to conclude that there is therefore no slowdown associated with the low TFP growth sector taking over the entire economy. However, consistent with the message from Figure 7, Duernecker et al. (2017a) show that if one measures output growth in the model through chained Fisher indexes as it is done by the BEA, then one would observe a decrease in measured output growth over time along the ABGP.

Lastly, we note that Acemoglu and Guerrieri (2008) generated a result similar to our Proposition 4. Differently from us, they imposed that investment and consumption have identical production functions, implying that investment necessarily mimics consumption, and hence that the overall economy also necessarily mimics consumption. Our analysis shows that this result occurs along an ABGP also when one allows for different production functions in consumption and investment and there is exogenous investment-specific technical change. Specifically, our result holds without imposing any restrictions on the behavior of A_x relative to the other TFPs, other than that A_{xt} must not shrink.

29. Related work by Duernecker et al. (2017b) disaggregates services and finds that the services subsector with the slowest productivity growth may not take over the economy asymptotically. The reason for this is that in the data services with fast and slow productivity growth appear to be substitutes, instead of complements.

9. IS INVESTMENT-BIASED TECHNICAL CHANGE ENDOGENOUS?

In this subsection, we present a third insight from modelling structural change in investment, namely, that technical change is endogenously investment-biased. Since our model with log-CES utility has a two-sector representation, one might think that technical change at the level of investment and consumption must be purely exogenous. Instead, the *effective* TFPs of the pseudo production functions in our two-sector representation are non-linear functions of the three *primitive* TFPs. Modelling structural change in investment is therefore essential if one seeks to draw inference about the nature of investment-specific technical change.

9.1. *Theoretical Analysis of Investment-biased Technical Change*

To see what we mean by technical change being *endogenously* investment biased, it is useful to recall the alternative production structure (5.6)–(5.7) that abstracts from structural change in investment. In the equilibrium of that framework, we have:

$$P_{jt} = \frac{B_{xt}}{B_{jt}}, \qquad j \in \{g, s\}.$$

It is common in the literature to measure investment-biased technical change by focusing on the price of nondurable consumption relative to equipment. Given our level of disaggregation, the most natural model counterpart to this is the price of services relative to investment, which, given our normalization of $P_{xt} = 1$, is simply equal to P_{st}. In the above specification, the sectoral TFPs B_{jt} are viewed as exogenous and one can infer the growth in B_{xt} relative to B_{st} simply by observing the change in the relative price of services. In other words, the behavior of this relative price is sufficient to infer the extent of investment-biased technical change.

This is in contrast to the inference in our model with structural change in investment. The expression for $\mathcal{A}_{xt}$ that we derived previously implies that relative prices in terms of primitive TFPs are given by:

$$P_{jt} = \frac{\mathcal{A}_{xt}}{A_{jt}} = \frac{A_{xt}\left(\omega_x A_{gt}^{\varepsilon_x - 1} + (1-\omega_x) A_{st}^{\varepsilon_x - 1}\right)^{\frac{1}{\varepsilon_x - 1}}}{A_{jt}}, \qquad j \in \{g, s\}.$$

The key message from these two expressions is that in each case the price of consumption relative to investment reflects an element of *endogenous* investment-biased technical change. Focusing on the relative price of services, it remains true that P_{st} increases if A_{xt} increases more strongly than A_{st} holding all else constant. But, it is also true that P_{st} will increase even absent any change in A_{xt} if A_{gt} increases more strongly than A_{st}. This effect reflects the *endogenous* input choices by the representative firm producing investment. That is, increases in A_{xt} lead to higher TFP in the production of investment independently of how the firm chooses to produce investment, but changes in A_{gt} relative to A_{st} influence the effective TFP of investment to an extent that depends on how the firm chooses to produce investment. Because the literature has identified investment-biased technical change with the *exogenous* movements in B_{xt}, we refer to this additional channel in our model as *endogenous* investment-biased technical change. We note that this occurs along any equilibrium path and is not restricted an ABGP. Formally, we summarize this discussion with the following proposition:

Proposition 5. *If $\widehat{A}_{gt} > \widehat{A}_{st}$, and $\widehat{A}_x = 0$, then along any equilibrium path $\widehat{P}_{st} > 0$.*

In sum, the key point of the previous discussion is that the price of services relative to investment will increase even in the absence of any exogenous investment-specific technical change. If $\widehat{A}_{xt} > 0$ then there will be an additional source of growth in P_{st}.

In the context of our model, an alternative to studying the price of services relative to investment is to study the price of consumption relative to investment. This is possible with log-CES utility when our model reduces to a two-sector model with consumption and investment that has an analytic expression for the price of consumption relative to investment.[30] To see this, note that given the symmetry of the aggregators in the investment and consumption sectors in the log-CES case, applying the earlier derivations that we made for the investment sector to the consumption sector with log-CES utility, we find that:

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

where

$$\mathcal{A}_{ct} \equiv \left(\omega_c A_{gt}^{\varepsilon_c - 1} + (1-\omega_c) A_{st}^{\varepsilon_c - 1}\right)^{\frac{1}{\varepsilon_c - 1}},$$

$$L_{ct} \equiv \frac{C_{gt}}{A_{gt} K_t^{\theta}} + \frac{C_{st}}{A_{st} K_t^{\theta}}, \quad K_{ct} \equiv K_t L_{ct}.$$

This implies that

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

In Appendix C below, we prove the following result:

Proposition 6. *Suppose the period-utility function is log-CES. If $\omega_x > \omega_c$, $\varepsilon_x \geq \varepsilon_c$, $\widehat{A}_{gt} > \widehat{A}_{st}$, and $\widehat{A}_x = 0$, then along any equilibrium path $\widehat{P}_{ct} > 0$.*

Once again, the key implication is that the price of consumption relative to investment will increase even in the absence of any exogenous investment-specific technical change.

9.2. *Empirical Analysis of Investment-biased Technical Change*

The preceding subsection argued theoretically that our model featuring structural change in the investment sector contains a novel channel that influences the relative price of consumption to investment. In this subsection we evaluate the empirical significance of this channel.

Figure 6 above established that the growth rates of both A_{st} and A_{xt} are relatively small when compared to the growth rate of A_{gt}. This suggests that the endogenous component of investment-biased technical change will play a large role in accounting for the behavior of the price of services relative to investment, P_{st}. The left part of Figure 8 provides information on this, showing the behavior of A_{xt}/A_{st}, which reflects

30. Unlike the log-CES case, the PIGL case does not permit one to reduce our model to the standard two-sector model with investment and consumption and to derive an analytic expression for the price of aggregate consumption relative to aggregate investment.

FIGURE 8
The components of investment-specific technical progress

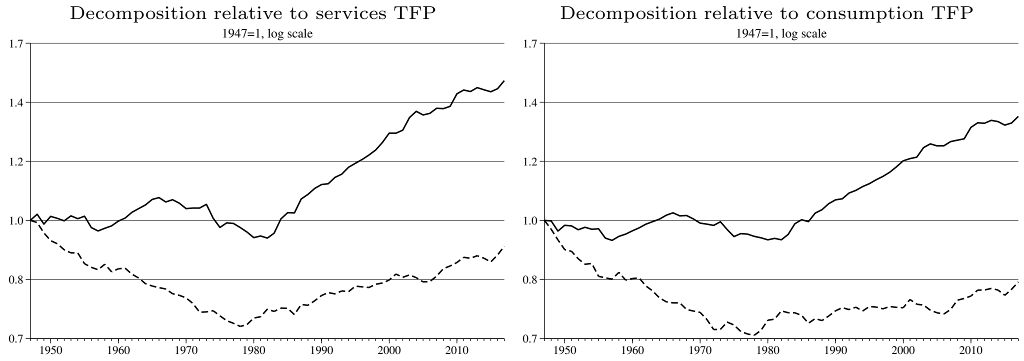

the contribution of exogenous investment-biased technical change to P_{st}, and $\mathcal{A}_{xt}/A_{st}$, which reflects the model implied value of P_{st}. The key message to take away from these curves is that exogenous investment biased technical change has not been a quantitatively important source of increases in this relative price in the US. Even focusing on the post-1980 period we see that it accounts for only about one quarter of the overall increase in the price of services relative to investment. Garcia-Santana et al. (2018) reach a similar conclusion in their analysis of cross-country patterns: exogenous investment-biased technical change accounts for only a third of the decline in the relative price of investment; the rest is accounted for by endogenous investment-biased technical change.

If we focus on the log-CES utility specification then we can also compute a model based price for aggregate consumption relative to investment, P_{ct}. The right-hand panel of Figure 8 displays these results. In this figure, $\mathcal{A}_{xt}/\mathcal{A}_{ct}$ reflects the model implied price of consumption relative to investment, and $A_{xt}/\mathcal{A}_{ct}$ reflects the contribution of exogenous investment-biased technical change to the price of consumption relative to investment. The take-away message remains the same as for the previous figure: in the post-1980 period, exogenous investment-biased technical change accounts for only about one quarter of the increase in P_{ct}.

10. CONCLUSION

In this paper, we have proposed a new framework for studying growth and structural change in both investment and consumption. Production has taken place at the level of sector value added and final expenditures on investment and consumption have been aggregates of the underlying sectoral value added. We have shown that this framework has three distinctive implications. First, constant growth in all sectoral TFPs is generically inconsistent with aggregate balanced growth and structural change in investment. Second, the sector with the lowest productivity growth asymptotically dominates the entire economy. Third, for empirically plausible parameter values technical change is endogenously investment specific. In addition, we have shown that 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, it will be of interest to pursue more fully a formulation that separates investment into the three components: structures, equipment, and IPP. We believe that this framework will help us better isolate the underlying sources of TFP growth.[31] It will also be of interest to disaggregate further the goods and services that are used to produce investment and consumption. One obvious example for why this can matter is that productivity growth differs significantly across sectors within the services sector, and the investment sector may use a different mix of services than the consumption sector.

Lastly, we have focused on the post-war US. While that is a natural starting point, we think that it will be interesting to extent our analysis to other developed countries and evaluate to what extent the necessary condition for ABGP – that $\widehat{\mathcal{A}}_{xt}$ is constant on average – holds more broadly. Doing this will become feasible once reasonably long time series of input-output tables become available also for developed countries other than the U.S.

A. PROOF OF LEMMA 1

We start by rewriting the production function of X_t in two different ways:

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

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

Equation (5.5) implies that:

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

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

Combining the last two sets of equations, we obtain:

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

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

Rearranging gives:

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

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

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 both sides by 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**

31. In the working-paper version of this paper, Herrendorf et al. (2018), we presented some preliminary evidence for the decomposition of investment into structures, equipment, and intellectual property products.

B. PROOF OF PROPOSITION 1

Part 1 – Necessity. We start by showing that if an ABGP exists, then $\widehat{\mathcal{A}}_{xt}$ is constant and (6.14) holds.

Since Y_t and K_t are growing at constant rates along an ABGP, the aggregate production function, $Y_t = \mathcal{A}_{xt} K_t^\theta$, implies that $\mathcal{A}_{xt}$ also grows at a constant rate.

This leaves to show that all aggregate variables grow at a common rate and R is constant. The Euler–equation can be written as

$$-\frac{\dot{V}_E(t)}{V_E(t)} = (1-\chi)\widehat{E}_t + \chi\left(\widehat{\mathcal{A}}_{xt} - \widehat{A}_{st}\right) = R_t - \delta - \rho, \tag{B22}$$

where we used the fact that $P_{st} = \mathcal{A}_{xt}/A_{st}$. $\widehat{E}_t$ is constant along an ABGP and we have shown in the previous step that $\widehat{\mathcal{A}}_{xt}$ is constant. Thus, our assumption that $\chi\widehat{A}_{st}$ is constant implies that the right-hand side is constant, and so the rental price of capital is constant, $R_t = R$. Since

$$\theta \mathcal{A}_{xt} K_t^{\theta-1} = R_t = R, \tag{B23}$$

it must be that $\widehat{K} = \widehat{\mathcal{A}}_x/(1-\theta)$. Then the aggregate production function, $Y_t = \mathcal{A}_{xt}K_t^\theta$, implies that $\widehat{Y} = \widehat{K}$. Next, to show that $\widehat{X} = \widehat{K}$, rewrite the capital accumulation equation as:

$$\frac{X_t}{K_t} = \widehat{K}_t + \delta.$$

Since $\widehat{K}_t$ is constant, X_t/K_t must be constant too. Thus, $\widehat{X} = \widehat{K}$.

Since $\widehat{E}_t$, $\widehat{K}_t$, $\widehat{X}_t$, and $\widehat{Y}_t$ must be constant and we have shown that $\widehat{K} = \widehat{X} = \widehat{Y}$, it remains to be shown that along an ABGP $\widehat{K} = \widehat{E}$. Combining $\dot{K}_t = X_t - \delta K_t$ with the resource constraint $X_t + E_t = Y_t$ gives:

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

$\widehat{K}$ is constant. Since $\widehat{K} = \widehat{Y}$, it follows that Y_t/K_t is constant. Hence, E_t/K_t must also be constant, implying that $\widehat{K} = \widehat{E}$.

In sum, we have shown that if an ABGP exists, then $\widehat{\mathcal{A}}_{xt}$ is constant and $\widehat{K} = \widehat{X} = \widehat{Y} = \widehat{E} = \widehat{\mathcal{A}}_x/(1-\theta)$. Thus, (6.14) holds.

Part 2 – Sufficiency. We continue by showing that if $\widehat{\mathcal{A}}_{xt} = \widehat{\mathcal{A}}_x$ is constant, then an ABGP exists. We do so by constructing one. Set $\widehat{K} = \widehat{X} = \widehat{Y} = \widehat{E} = \widehat{\mathcal{A}}_x/(1-\theta)$. Moreover, set R such that the Euler equation holds:

$$-\frac{\dot{V}_E(t)}{V_E(t)} = (1-\chi)\frac{\widehat{\mathcal{A}}_x}{1-\theta} + \chi\left(\widehat{\mathcal{A}}_x - \widehat{A}_s\right) = R - \delta - \rho.$$

Condition (6.13) in Assumption 2 together with Assumption 3 imply that there exists a constant $R > 0$ such that the equation holds. Then there is also a unique value of K_0 such that Equation (B23) is satisfied for this R and the given $\mathcal{A}_{x0}$. Next, pick X_0 such that the capital-accumulation equation is satisfied in instant 0 given K_0 and $\widehat{K}$. Then substitute the capital-accumulation equation into the feasibility constraint and pick E_0 such that the resulting equation is satisfied in instant 0 given K_0, $\mathcal{A}_{x0}$, and $\widehat{K}$:

$$E_0 = \mathcal{A}_{x0}K_0^\theta - \left[\widehat{K} + (1-\delta)\right] K_0.$$

Since all variables in the capital-accumulation equation and the feasibility constraint grow at the same constant rate, the two will be satisfied in every instant.

Lastly, we have to show that the transversality condition holds along the ABGP just constructed. This amounts to showing:

$$\lim_{t\to\infty} e^{-\rho t} E_t^{\chi-1} P_{st}^{-\chi} K_t = \lim_{t\to\infty} e^{-\rho t}\left(\frac{E_t}{P_{st}}\right)^\chi \frac{K_t}{E_t} = 0.$$

Given that along an ABGP $\widehat{K} = \widehat{E}$, this condition holds if and only if:

$$-\rho + \chi\frac{\widehat{\mathcal{A}}_x}{1-\theta} - \chi\left(\widehat{\mathcal{A}}_x - \widehat{A}_s\right) < 0.$$

This is a version of Condition (6.12) in Assumption 2. **QED**

C. PROOF OF PROPOSITION 6

We start with the well-known result that for log-CES instantaneous utility the model implied P_{ct} equals the Divisia index that the BEA uses in NIPA. Given that consumption and investment are composites of goods and services in the model, the Divisia index for the change in the price of consumption relative to investment is given by:

$$\widehat{P}_{ct} = \left(\zeta^c_{gt}\widehat{P}_{gt} + (1-\zeta^c_{gt})\widehat{P}_{st}\right) - \left(\zeta^x_{gt}\widehat{P}_{gt} + (1-\zeta^x_{gt})\widehat{P}_{st}\right) = (\zeta^c_{gt} - \zeta^x_{gt})\left(\widehat{P}_{gt} - \widehat{P}_{st}\right), \tag{C24}$$

where

$$\zeta^c_{gt} \equiv \frac{P_{gt}C_{gt}}{P_{gt}C_{gt} + P_{st}C_{st}},$$
$$\zeta^x_{gt} \equiv \frac{P_{gt}X_{gt}}{P_{gt}X_{gt} + P_{st}X_{st}}.$$

The product after the last equality sign in (C24) features two components: the difference in the goods shares in consumption expenditure versus investment expenditure; the growth rate of the price of goods relative to services. We want to show that both terms of the right-hand side are negative, implying that the right-hand side is positive.

The second terms on the right-hand side is negative because:

$$\frac{P_{gt}}{P_{st}} = \frac{A_{st}}{A_{gt}} \quad \Longrightarrow \quad \widehat{P}_{gt} - \widehat{P}_{st} = \widehat{A}_{st} - \widehat{A}_{gt} < 0.$$

The inequality follows from the assumption that $\widehat{A}_{gt} > \widehat{A}_{st}$

The first term of the right-hand side is negative if $\zeta^x_{gt} > \zeta^c_{gt}$. To see why this is the case, note that (5.5) and (6.16) imply that:

$$\zeta^x_{gt} = \frac{1}{1 + \dfrac{1-\omega_x}{\omega_x}\left(\dfrac{A_{gt}}{A_{st}}\right)^{1-\varepsilon_x}},$$
$$\zeta^c_{gt} = \frac{1}{1 + \dfrac{1-\omega_c}{\omega_c}\left(\dfrac{A_{gt}}{A_{st}}\right)^{1-\varepsilon_c}}.$$

A sufficient condition for $\zeta^x_{gt} > \zeta^c_{gt}$ is that $\omega_x > \omega_c$, $\varepsilon_x \geq \varepsilon_c$, and $A_{gt} \geq A_{st}$. Since our normalization $p_{s0}/p_{g0} = 1$ is equivalent to $A_{g0} = A_{s0} = 1$, the assumption that $\widehat{A}_{gt} > \widehat{A}_{st}$ ensures that $A_{gt} \geq A_{st}$. **QED**

Acknowledgements. We thank the editor Dirk Krueger, four anonymous referees, Timo Boppart, Julieta Caunedo, Per Krusell, Rachel Ngai, B. Ravikumar, Todd Schoellman, V. Kerry Smith, and the participants of the ASU Macro Workshop, the Conference on Growth and Productivity at the University of Toronto, the Cornell–Notre Dame Macro-Development Workshop, the Cowles Conference on Macro, the London School of Economics, the NBER Growth Group in San Francisco, Queen Mary University, the SED Meetings at ITAM, the Universities of Groningen, Köln, Lancaster, Western Ontario, and Windsor, the Vienna Macroeconomics Workshop, and the Workshop on Structural Transformation and Macroeconomic Dynamics in Cagliari for comments and suggestions. Valentinyi thanks the Hungarian National Research, Development and Innovation Office (Project KJS K 124808). All errors are our own.

REFERENCES

Acemoglu, D. and V. Guerrieri (2008). Capital deepening and non-balanced economic growth. *Journal of Political Economy 116*, 467–498.

Alvarez-Cuadrado, F., N. V. Long, and M. Poschke (2017). Capital–labor substitution, structural change, and growth. *Theoretical Economics 12*, 1229–1266.

Baumol, W. J. (1967). Macroeconomics of unbalanced growth: The anatomy of the urban crisis. *American Economic Review 57*, 415–426.

Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica 82*, 2167–2196.

Buera, F. J. and J. P. Kaboski (2009). Can traditional theories of structural change fit the data? *Journal of the European Economic Association 7*, 469–477.

Comin, D. and B. Hobijn (2010). An exploration of technology diffusion. *American Economic Review 100*, 2031–2059.

Comin, D., M. Mestieri, and D. Lashkari (2018). Structural change with long-run income and price effects. Manuscript.

Duernecker, G., B. Herrendorf, and Á. Valentinyi (2017a). Quantity measurement, balanced growth, and welfare in multi–sector growth models. Discussion Paper 12300, Centre for Economic Policy Research, London.

Duernecker, G., B. Herrendorf, and Á. Valentinyi (2017b). Structural change within the service sector and the future of baumol's cost disease. Discussion Paper 12467, Centre for Economic Policy Research, London.

Echevarria, C. (1997). Changes in sectoral composition associated with economic growth. *International Economic Review 38*, 431–452.

Garcia-Santana, M., J. Pijoan-Mas, and L. Villacorta (2018). Investment demand and structural change. Manuscript.

Herrendorf, B., C. Herrington, and Á. Valentinyi (2015). Sectoral technology and structural transformation. *American Economic Journal: Macroeconomics 7*, 1–31.

Herrendorf, B., R. Rogerson, and Á. Valentinyi (2013). Two perspectives on preferences and structural transformation. *American Economic Review 103*, 2752–2789.

Herrendorf, B., R. Rogerson, and Á. Valentinyi (2014). Growth and structural transformation. In P. Aghion and S. N. Durlauf (Eds.), *Handbook of Economic Growth*, Volume 2, pp. 855–941. Elsevier.

Herrendorf, B., R. Rogerson, and Á. Valentinyi (2018). Structural change in investment and consumption: A unified approach. Working Paper 24568, National Bureau of Economic Research, Cambridge, MA.

Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and manufacturing tfp in china and india. *Quarterley Journal of Economics 74*, 1403–1448.

Kongsamut, P., S. Rebelo, and D. Xie (2001). Beyond balanced growth. *Review of Economic Studies 68*, 869–882.

Ngai, L. R. and C. A. Pissarides (2007). Structural change in a multisector model of growth. *American Economic Review 97*, 429–443.

Solow, R. (1957). Technical change and the aggregate production function. *Review of Economics and Statistics 38*, 312–320.

Sposi, M. (2012). Evolving comparative advantage, structural change, and the composition of trade. Jobmarket paper, University of Iowa.