The Past and Future of U.S. Structural Change: Compositional Accounting and Forecasting*

Andrew T. Foerster
Federal Reserve Bank of San Francisco

Andreas Hornstein
Federal Reserve Bank of Richmond

Pierre-Daniel G. Sarte
Federal Reserve Bank of Richmond

Mark W. Watson
Princeton University and NBER

January 27, 2026

Abstract

We explore the evolving significance of different production sectors within the U.S. economy since World War II and provide methods for estimating and forecasting these shifts. Using a compositional accounting approach, we find that the well-documented transition from goods to services is primarily driven by two compositional changes: 1) the rise of Intellectual Property Products (IPP) as an input producer, replacing Durable Goods almost one-for-one in terms of input shares in virtually all sectors; and 2) a shift in consumer spending from Nondurable Goods to Services. A structural model replicating these shifts reveals that the rise of IPP at the expense of Durable Goods is largely explained by increases in the efficiency of IPP inputs used in production: input-biased technical change. Trend variations in sectoral total factor productivity, and their attendant effects on relative prices and income, are the main driver of evolving consumption patterns. Both reduced-form and structural forecasts project these trends to continue over the next two decades, albeit at lower rates, indicating a slower pace of structural change.

Keywords: Structural Change, Compositional Accounting, Structural Forecasting, Input Biased Technical Change, Low Frequency Trends

JEL Codes: E17, E23, E27

*The views expressed in this paper are those of the authors and do not necessarily represent the views of the Federal Reserve Banks of Richmond or San Francisco, nor those of the Federal Reserve System. We thank Christian vom Lehn for thoughtful insights, as well as participants at various seminars and conferences for helpful comments. Finally, we thank Erin Henry, Conner Mulloy and Jack Taylor for outstanding research assistance.

Andrew Foerster: andrew.foerster@sf.frb.org, Andreas Hornstein: and.horn.re@gmail.com, Pierre-Daniel Sarte: pierre.sarte@rich.frb.org, Mark Watson: mwatson@princeton.edu.

1 Introduction

Since the end of WWII, the sectoral composition of U.S. GDP is well known to have shifted markedly from goods to services.[1] Underlying this general shift, however, are other underlying secular compositional changes that give a more detailed picture of structural change. This paper explores these compositional trends both from a historical accounting standpoint – where the evolution of sectoral value-added shares of GDP reflect changes in constituent subshares, for example in consumption expenditure shares or in intermediate input cost shares – and from the standpoint of an economic model – where all compositional shifts are driven by more fundamental forces such as productivity or intrinsic preferences for certain goods. We further introduce methods for forecasting compositional trends both in reduced-form, within our compositional accounting framework, and in structural form, within the model where the composition of the U.S. economy evolves endogenously in general equilibrium.

We begin with compositional accounting, and the notion that sectoral value-added shares of GDP ultimately reflect sets of constituent subshares that describe the various ways in which a sector's gross output might be used, all measured relative to its value added. Thus, sectoral resource constraints dictate that each sector's share of GDP must reflect a mix of consumption expenditure shares, either private or public, input cost shares, either in materials or investment, and trade shares. Moreover, because these constituent subshares sum to unity, they do not evolve independently of each other either within or across sectors.

A decomposition of sectoral shares of GDP over the post WWII period reveals that the steady shift from goods to services was primarily driven by two key compositional changes. First, Intellectual Property Products (IPP) gradually displaced Durable Goods as a source of inputs in virtually all production sectors, both as materials and investment, almost one-for-one in terms of input cost shares. As a share of GDP, IPP stands today approximately where Durable Goods stood in 1947. Second, the consumption share of Services steadily increased over the postwar period in a way that was almost exactly offset by a decline in the consumption share of Nondurable Goods, essentially leaving their combined consumption shares unchanged throughout. More generally, while goods versus services is a useful taxonomy of sectors, it is also helpful to classify sectors in terms of those that contribute mainly to inputs for production versus those that contribute mainly to final consumption. Viewed in this way, we show that the GDP share of input-producing sectors has stayed remarkably constant throughout the whole post WWII period and, therefore, so has that of consumption-producing sectors.

Compositional accounting is helpful for breaking down the evolution of sectoral shares of GDP into underlying secular shifts in consumption or production. However, it cannot tell us how these shifts are related to more fundamental forces, such as changes in sectoral productivity or preferences towards consumption from particular sectors. To that end, we construct a structural model that captures production linkages in intermediate goods and investment, income effects and relative price

[1] For a survey of structural change over time and across countries see Herrendorf, Rogerson, and Valentinyi (2014) or Restuccia and Rogerson (2017).

effects on consumption, and labor supply shifts. We then use this model to quantify how different forces, both on the supply side and on the demand side, have contributed to structural change in general equilibrium. Exogenous supply-side forces are associated with changes in sector-specific total factor productivity (TFP) as well as changes in the relative efficiency with which inputs translate into output, defined here as input-biased technical change (IBTC). Changes on the demand side arise endogenously from income effects associated with overall productivity increases, from exogenous changes in intrinsic preferences for specific goods, and from changes in net exports and government spending.

We find that the gradual shift in post WWII consumption from Nondurables to Services is explained almost equally by income effects resulting from a general increase in productivity, and relative price effects arising from changes in relative sectoral productivity. I2606nterestingly, exogenous changes in preferences play a generally minor role in explaining this shift, and essentially no role prior to the 1980s. The gradually increasing share of IPP in GDP, and its rise as an input replacing Durable Goods, stems primarily from input-biased technical change, in particular relative increases in the efficiency of IPP inputs used in production across sectors. Together, these forces explain the bulk of U.S. structural change since WWII.

Beyond a structural historical account of compositional changes, we tackle the question of how to forecast future U.S. structural change. We do this in two ways. First, we apply versions of the long-run prediction methods developed by Müller and Watson (2016) to the problem of reduced-form compositional forecasting. Because shares lie between zero and one and sum to unity, this problem involves a transformation of the data that enforces these constraints but that is not necessarily motivated by economic analysis.[2] We then present a second approach in which trend compositional forecasts are constructed from long-run forecasts of fundamentals underlying our structural model. Under the latter approach, preferences, technologies, and resource constraints enforce the constraints on sectoral shares. This approach also allows us to examine the extent to which fundamental forces such as productivity change affect the future composition of the economy. Both reduced-form and structural forecasts project historical trends to continue over the next two decades, albeit at reduced rates, indicating a slowdown in the pace of structural change.

Our analysis contributes to a series of related papers starting with Ngai and Pissarides (2007) in which, given Constant Elasticity of Substitution (CES) preferences, sectoral consumption expenditure shares vary in response to changes in relative prices driven by sector-specific productivity. Herrendorf, Rogerson, and Valentinyi (2013) then underscore the distinction between consumption and value added when considering the role of income effects on the consumption of services for the transition from goods to services. Beyond consumption, Herrendorf, Rogerson, and Valentinyi (2021) further emphasize the significance of services in the production of investment in this transition. Closer to the production network literature, Gaggl, Gorry, and Vom Lehn (2023) highlight the significance of changes in both input-output (IO) and capital-flow linkages in explaining the compositional shifts towards services.

We share with García-Santana, Pijoan-Mas, and Villacorta (2021) a focus on investment, and

[2] See the treatment in Aitchison (1986).

in our case the rising role of IPP as a source of new capital, in explaining structural change. However, unlike García-Santana et al. (2021), our analysis remains closer in spirit to Boppart, Kiernan, Krusell, and Malmberg (2023) in modeling this change as a sequence of steady states. Interestingly, García-Santana et al. (2021) estimate that goods and services are complements in the production of investment while Gaggl et al. (2023) estimate that they are substitutes. We find that Durable Goods and IPP can be either complements or substitutes in the production of both investment and materials depending on the sector employing these inputs.

Our analysis retains the insights from this work but also departs from it in important ways. First, while Herrendorf et al. (2021) and Gaggl et al. (2023) rely on data transformations to be consistent with a broad model-specific taxonomy of goods versus services, we frame our analysis using standard National Income Accounting industry definitions. Doing so helps us highlight that within goods and services, different types of goods, e.g., Durables versus Nondurables, and different types of services, e.g., IPP versus other services, have contributed to structural change in very different ways. The emphasis on Durables as a key producer of investment goods follows in the vein of Greenwood, Hercowitz, and Krusell (1997) and the subsequent literature on investment-specific technical change. The focus on IPP allows us to address the emerging role of the 'knowledge sector' as a producer of investment goods. Put simply, not all goods and not all services serve the same purpose, and hence play different roles in shaping the composition of the U.S. economy. Second, standard industry definitions make it possible to use information on sources of final demand and intermediate input use from the Bureau of Economic Analysis (BEA) to parameterize how an industry's gross output contributes to final demand and intermediate inputs. Our model, therefore, is closely aligned with the observed heterogeneity across sectors in capturing important drivers of structural change.[3] In addition, the model analysis can then be informed by corresponding measures of industry productivity obtained from KLEMS data. Finally, we do not impose common technologies in the production of gross output, value added, or investment goods across industries, which is typically motivated by balanced growth considerations.

This paper is organized as follows. Section 2 introduces the data and describes the long-run trends that form the basis of our empirical analysis. Section 3 introduces the notion of compositional accounting. It decomposes the evolution of sectoral value added shares into constituent subshares that describe the different uses of sectoral output. This section also carries out and presents reduced-form forecasts of trend sectoral shares of GDP. Section 4 lays out a model in which the composition of the economy responds to economic fundamentals and slowly evolves over time. Section 5 quantifies the model and establishes some benchmark results. Section 6 explores the role of various economic fundamentals in driving structural change and produces compositional forecasts constructed within our structural model. Section 7 concludes. A detailed online Technical Appendix contains a comprehensive description of the data, of the economic environment as well as all equilibrium conditions and solution algorithm, of the statistical methods used for estimation and long-run forecasting, and includes additional figures referenced in the main text.

[3]Put differently, production networks are an integral part of the model rather than subsumed in the construction of the data to be consistent with the model.

2 Data and Trends

In this section, we introduce the data used in our analysis, discuss changes in the composition of U.S. GDP since WWII, and give an overview of the computation of low-frequency trends that underpin our empirical results.

2.1 Data

Our analysis relies on industry-level data related to productivity, intermediate input use, sources of final demand and investment flows. We obtain these data from three sources. First, information on industry productivity is jointly provided by the Bureau of Economic Analysis (BEA) and the Bureau of Labor Statistics (BLS) through their Integrated Industry-Level Production Accounts (KLEMS). Second, information regarding intermediate input use and sources of final demand are obtained from the BEA's Input-Output tables.[4] Finally, data on investment flows are constructed and made available by vom Lehn and Winberry (2021) (vLW).[5] These datasets are attractive for our purposes because they provide a unified approach to production, consumption, and investment at the industry level. The KLEMS data build on U.S. National Income and Product Accounts (NIPA) and so consistently integrate industry data with Input-Output tables and Fixed Asset tables. The vLW investment flows provide additional details on sources of industry investment.

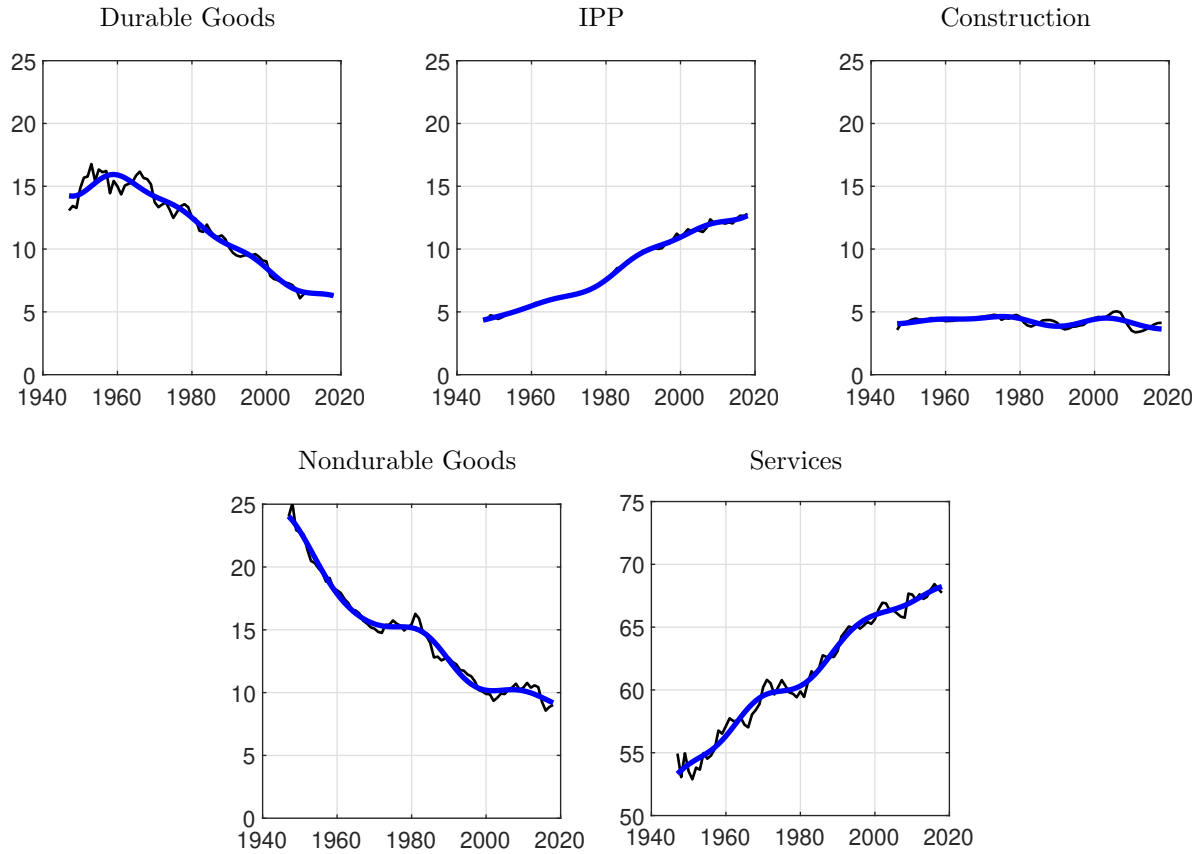
Our three datasets provide annual observations from 1947 to 2018. We use the KLEMS data to construct price and quantity indices at each date, t, and for each industry, j. These quantity indices and prices include gross output, denoted by $(y_{j,t}, p^y_{j,t})$, value-added, $(v_{j,t}, p^v_{j,t})$, intermediate aggregates, $(m_{j,t}, p^m_{j,t})$, capital services, $(k_{j,t}, u_{j,t})$, and labor services, $(\ell_{j,t}, w_{j,t})$. Capital income shares are denoted by $\alpha_{j,t} = (u_{j,t}k_{j,t})/(p^v_{j,t}v_{j,t})$. Industry total factor productivity (TFP), denoted by $z_{j,t}$, is defined as the Solow residual in each industry − that is, industry value added less the services of capital and labor. We use the Input-Output tables to calculate pairwise intermediate input shares across all industries i and j, shares of private consumption by industry, and ratios of government consumption and net-exports to value-added for each industry. Data on investment flows give us pairwise investment cost shares across all industries i and j.

We consolidate our industry data into five sectors that produce structures, durable goods, intellectual property products (IPP), nondurable goods, and services. We choose this particular breakdown in part because it helps underscore the key forces and mechanisms that have affected the compositional evolution of the U.S. economy. Aggregate quantity indices for the different sectors, for example combining all durable goods industries into a single Durable Goods sector, are constructed as Divisia indices of the underlying industry series.

[4] For a discussion of how the input-output network has changed over time, see Choi and Foerster (2017).

[5] Casal and Caunedo (2024) construct similar measures from different countries and compare how the investment network changes with stages of development.

Figure 1: Trends in the Sectoral Composition of U.S. GDP
1947-2018 (Percentage Points)



Notes: Each panel shows the share of nominal GDP associated with each sector. Thin lines are the raw data and thick lines are low-frequency trends.

2.2 The Changing Composition of GDP

Figure 1 plots the shares of GDP associated with each of the five sectors. The figure shows the raw data along with an associated low-frequency trend (described below) that captures secular changes in the economy. Two observations stand out.

First, the well-documented shift from goods to services is apparent. In the context of our five key sectors, goods comprise Durable Goods, such as motor vehicles, industrial machinery, and consumer electronics, Nondurable Goods such as food and apparel, and Construction. Services broadly defined include Intellectual Property Products (IPP), mainly composed of software, research & development, as well as technical, scientific and professional services, and other Services such as transportation, leisure & hospitality, and health and education. We will refer to these two sectors simply as IPP and Services and use the term 'Total Services' when we need to denote their sum. Over our sample's 72-year time span, the share of goods (Construction+Durables+Nondurables) in GDP decreased by more than 20 percentage points, offset by a corresponding increase in Total Services (IPP+Services).

Second, as indicated in Table 1, the five-sector decomposition provides a more nuanced view of structural change. In particular, Services and Nondurable Goods (which are primarily used for

Table 1: Changes in Sectoral Value Added Shares
1947 to 2018 (Percentage Points)

Sectors	Durable Goods	IPP	Construction	Nondurable Goods	Services
Trend Value in 1947	14.2	4.3	4.1	24.0	53.3
Trend Value in 2018	6.3	12.7	3.6	9.2	68.2
Change in Trend Value	-8.0	8.3	-0.4	-14.9	14.9

consumption) show large offsetting changes of 15 percentage points, while IPP and Durable Goods (which are primarily used as inputs to production) show similarly offsetting changes of 8 percentage points; the share of Construction remained approximately constant throughout the sample period.

This paper is concerned with explaining these historical structural shifts and forecasting future shifts. Figure 1 makes clear that the analysis of structural change is inherently a low-frequency problem. Thus, the trends plotted in the figure, together with trends of other variables, play a key role in our analysis. Before starting that analysis, we take a detour to describe how these trends are computed.

2.3 Low-frequency Trends

Let x_t denote a time series, and let $\hat{x}_t$ denote its long-run trend. We compute $\hat{x}_t$ using a 'low-pass' filtering method described in a series of papers, specifically Müller and Watson (2008, 2016, 2017, 2018, 2020), and that we have used in previous work, Foerster, Hornstein, Sarte, and Watson (2022). In particular, $\hat{x}_t$ is computed as the fitted value from the regression of x_t onto a constant, possibly a linear time trend, and a set of low-frequency periodic functions, using sample data from $t = 1, ..., T$. For regressions that include a constant but no time trend, the periodic functions are q cosines that jointly capture the variability in the time series for periods longer than $2T/q$. For regressions that include a time trend, the periodic functions are slightly different, but again $\hat{x}_t$ captures the same low-frequency periodicities. Appendix A provides a more detailed description of the process for computing $\hat{x}_t$.

In our dataset, $T = 72$ years and we set $q = 7$ so that $\hat{x}_t$ captures periods longer than $2 \times 72/7 \approx 20$ years. Thus, our interest is in the long-run evolution of the economy, where 'long-run' is *defined* as periodicities longer than 20 years. We modify this method when computing trends for shares, such as the series plotted in Figure 1. In particular, shares take on values between zero and one and sum to unity across sectors, and we impose this constraint on the trends. Following Aitchison (1986), we implement the constraint using a logit transformation, denoted L. Thus, let $s_{i,t}$ denote one of n shares with $s_{i,t} \in [0, 1]$ and $\sum_{i=1}^{n} s_{i,t} = 1$. We compute $x_{i,t} = \ln(s_{i,t}/s_{n,t})$, then the trends, $\hat{x}_{i,t}$, as described above, and form the share trends as $\hat{s}_{i,t} = \exp(\hat{x}_{i,t}) / \sum_{j=1}^{n} \exp(\hat{x}_{j,t})$. These are the trends plotted in Figure 1.

Given our interest in structural change, we use these trends as the source data for all of our calculations. For notational simplicity, we drop the '$\hat{}$' notation, but it should be assumed throughout.

Finally, unless stated otherwise, we include a linear time trend for computing the trends.

3 Compositional Accounting

Analogous to growth accounting, where resource constraints in the economy impose restrictions on the relationship between various growth rates, compositional accounting describes how different sectoral shares, such as sectoral shares of final consumption or input shares, interact to arrive at final sectoral shares of GDP. Compositional accounting, therefore, is helpful in describing historical patterns in the composition of the U.S. economy.

3.1 Historical Accounting: Trends in the Changing Composition of the U.S. Economy

To begin disentangling the sources of structural change underlying Figure 1, consider an economy with n production sectors, indexed by j (or i), collected in $\mathcal{N} = \{1, ..., n\}$. At date t, each sector j produces a final product, $y_{j,t}$, which can be purchased at price $p^y_{j,t}$ to be consumed, $c_{j,t}$, used as materials by another sector i, denoted $m_{ji,t}$, or used as investment by another sector i, denoted $x_{ji,t}$. Sector j goods can also be purchased by the government, $g_{j,t}$, or traded internationally, where $nx_{j,t}$ denotes net exports. We denote the real value added in sector j by $v_{j,t}$, with the corresponding nominal value added, $p^v_{j,t}v_{j,t}$.

Because different sectors purchase goods or services from other sectors in ways that change over time, the economy features an evolving network of sectoral linkages. Observations on the types of inputs that different sectors produce for other sectors, or on output they produce for final consumption, are available from the BEA Fixed Asset and Input-Output tables. Let capital letters denote nominal quantities, and lowercase letters real quantities. Then, for each sector j and each time period t, we observe the following.

- Expenditures on goods for consumption, $C_{j,t} = p^y_{j,t}c_{j,t}$, and as a share of total consumption expenditures, $\theta_{j,t} = C_{j,t}/C_t$, with $C_t = \sum_{j\in\mathcal{N}} p^y_{j,t}c_{j,t}$. Let $\boldsymbol{\theta}_t = (\theta_{1,t}, ..., \theta_{n,t})'$.

- Nominal value added, $V_{j,t} = p^v_{j,t}v_{j,t}$, and as a share of gross output, $\gamma_{j,t} = \frac{V_{j,t}}{Y_{j,t}} = \frac{p^v_{j,t}v_{j,t}}{p^y_{j,t}y_{j,t}}$, with $\boldsymbol{\gamma}_t = (\gamma_{1,t}, ..., \gamma_{n,t})'$ and $\boldsymbol{\Gamma}_t = diag\,(\gamma_{j,t})$. Nominal GDP is given by $\sum_{j\in\mathcal{N}} V_{j,t}$.

- Expenditures on investment goods, $X_{j,t} = p^x_{j,t}x_{j,t}$, and as a share of value added, $\psi^x_{j,t} = \frac{X_{j,t}}{V_{j,t}} = \frac{p^x_{j,t}x_{j,t}}{p^v_{j,t}v_{j,t}}$, with $\boldsymbol{\psi}^x_t = (\psi^x_{1,t}, ..., \psi^x_{n,t})'$ and $\boldsymbol{\Psi}^x_t = diag\left(\psi^x_{j,t}\right)$.

- Government expenditures, $G_{j,t} = p^y_{j,t}g_{j,t}$, and as a share of value-added, $\psi^g_{j,t} = \frac{G_{j,t}}{V_{j,t}} = \frac{p^y_{j,t}g_{j,t}}{p^v_{j,t}v_{j,t}}$, with $\boldsymbol{\psi}^g_t = (\psi^g_{1,t}, ..., \psi^g_{n,t})'$ and $\boldsymbol{\Psi}^g_t = diag\left(\psi^g_{j,t}\right)$.

- Net exports, $NX_{j,t} = p^y_{j,t}nx_{j,t}$, and as a share of value added, $\psi^{nx}_{j,t} = \frac{NX_{j,t}}{V_{j,t}} = \frac{p^y_{j,t}nx_{j,t}}{p^v_{j,t}v_{j,t}}$, with $\boldsymbol{\psi}^{nx}_t = (\psi^{nx}_{1,t}, ..., \psi^{nx}_{n,t})'$ and $\boldsymbol{\Psi}^{nx}_t = diag\left(\psi^{nx}_{j,t}\right)$.

- Expenditures by sector j on materials from other sectors, i, $M_{ij,t} = p^y_{i,t} m_{ij,t}$, as a share of j's total intermediate input expenditures, $\phi_{ij,t} = \frac{M_{ij,t}}{M_{j,t}} = \frac{p^y_{i,t} m_{ij,t}}{p^m_{j,t} m_{j,t}}$, where $p^m_{j,t} m_{j,t}$ denotes j's total expenditures on materials, with $\mathbf{\Phi}_t = [\phi_{ij,t}]$. By construction, the columns of $\mathbf{\Phi}_t$ sum to one, $\sum_{i \in \mathcal{N}} \phi_{ij,t} = 1$ for all j and t.

- Expenditures by sector j on investment from every other sector, i, $X_{ij,t} = p^y_{i,t} x_{ij,t}$, as a share of j's total investment expenditures, $\omega_{ij,t} = \frac{X_{ij,t}}{X_{j,t}} = \frac{p^y_{i,t} x_{ij,t}}{p^x_{j,t} x_{j,t}}$, where $p^x_{j,t} x_{j,t}$ denotes j's total investment expenditures, with $\mathbf{\Omega}_t = [\omega_{ij,t}]$. By construction, the columns of $\mathbf{\Omega}_t$ sum to one, $\sum_{i \in \mathcal{N}} \omega_{ij,t} = 1$ for all j and t.

These component shares summarize the different ways in which the output of a given sector is used.[6] Together, they determine the importance of the overall role of that sector in the economy, or its share of GDP, $s^v_{j,t} = \frac{V_{j,t}}{V_t}$, as depicted in Figure 1.

At each date t, each sector j satisfies the nominal resource constraint

$$p^y_{j,t} y_{j,t} = p^y_{j,t} c_{j,t} + \sum_{i \in \mathcal{N}} p^y_{j,t} m_{ji,t} + \sum_{i \in \mathcal{N}} p^y_{j,t} x_{ji,t} + p^y_{j,t} g_{j,t} + p^y_{j,t} nx_{j,t}, \tag{1}$$

which, given the notation introduced above, can be expressed relative to its nominal value added, $V_{j,t}$,

$$\frac{V_{j,t}}{\gamma_{j,t}} = \theta_{j,t} C_t + \sum_{i \in \mathcal{N}} \phi_{ji,t} (1 - \gamma_{i,t}) \frac{V_{i,t}}{\gamma_{i,t}} + \sum_{i \in \mathcal{N}} \omega_{ji,t} \psi^x_{i,t} V_{i,t} + \psi^g_{j,t} V_{j,t} + \psi^{nx}_{j,t} V_{j,t}.$$

Let $\boldsymbol{V}_t = (V_{1,t}, ..., V_{n,t})'$ denote the vector of sectoral nominal value added. Then, we can summarize all of the economy's nominal resource constraints in vector form,

$$\mathbf{\Gamma}_t^{-1} \boldsymbol{V}_t = \boldsymbol{\theta}_t C_t + \mathbf{\Phi}_t (\boldsymbol{I} - \mathbf{\Gamma}_t) \mathbf{\Gamma}_t^{-1} \boldsymbol{V}_t + \mathbf{\Omega}_t \mathbf{\Psi}^x_t \boldsymbol{V}_t + \mathbf{\Psi}^g_t \boldsymbol{V}_t + \mathbf{\Psi}^{nx}_t \boldsymbol{V}_t. \tag{2}$$

The vector of sectoral shares of GDP, denoted $\boldsymbol{s}^v_t = (s^v_{1,t}, ..., s^v_{n,t})'$ is then given by the *compositional identity*,

$$\boldsymbol{s}^v_t = \boldsymbol{\eta}_t / (\mathbf{1}' \boldsymbol{\eta}_t), \text{ where } \boldsymbol{\eta}_t = \frac{\boldsymbol{V}_t}{C_t} = \left([\boldsymbol{I} - \mathbf{\Phi}_t (\boldsymbol{I} - \mathbf{\Gamma}_t)] \mathbf{\Gamma}_t^{-1} - \mathbf{\Omega}_t \mathbf{\Psi}^x_t - \mathbf{\Psi}^g_t - \mathbf{\Psi}^{nx}_t \right)^{-1} \boldsymbol{\theta}_t. \tag{3}$$

We can summarize sector j's share of GDP in equation (3) as

$$s^v_{j,t} = \mathcal{S}_j \left(\boldsymbol{\theta}_t, \boldsymbol{\gamma}_t, \boldsymbol{\psi}^x_t, vec(\mathbf{\Phi}_t), vec(\mathbf{\Omega}_t), \boldsymbol{\psi}^g_t, \boldsymbol{\psi}^{nx}_t \right),$$

where it is evident that changes in $s^v_{j,t}$ arise in various ways, some related to compositional shifts in production, others to changes in consumption patterns, and others still to changes in international conditions via trade. In general, the compositional identity, $\mathcal{S}_j$, is a function of $2n^2 + 5n$ sectoral

[6] Annual nominal data on final goods expenditure shares, $\boldsymbol{\theta}_t$, intermediate goods shares, $\mathbf{\Phi}_t$, and sectoral value added shares in gross output, $\boldsymbol{\gamma}_t$, are obtained from the BEA Input Output tables, while capital flows, $\mathbf{\Omega}_t$, are obtained from vom Lehn and Winberry (2021), scaled to match industry investment from the BEA fixed asset tables. Gross output price indices are obtained from the KLEMS data.

subshares, $\theta_{i,t}$, $\gamma_{i,t}$, $\phi_{ij,t}$, etc. Therefore, in an environment with 5 sectors, changes in the relative aggregate importance of any sector, j, reflect 75 interrelated compositional changes. These compositional changes in turn summarize how the U.S. economy has evolved in the way that different goods and services are used to produce other goods and services, either as inputs or as final consumption: structural change.

The post-WWII evolution of the various shares depicted in our key compositional identity, (3), is summarized in Figure 2. Three variables are plotted in each panel of the figure. The value-added shares are plotted in black – these were shown previously in Figure 1. Consumptions shares are plotted in red; in the compositional identity (3), these correspond to value-added shares absent production linkages as well as government and net export demand, that is with $\boldsymbol{\Gamma}_t = \boldsymbol{I}$, $\boldsymbol{\Omega}_t = \mathbf{0}$, and $\boldsymbol{\Psi}_t^x = \boldsymbol{\Psi}_t^g = \boldsymbol{\Psi}_t^{nx} = \mathbf{0}$. The series plotted in blue additionally include production linkages (and so incorporates the historical values of $\boldsymbol{\theta}_t$, $\boldsymbol{\Gamma}_t$, $\boldsymbol{\Phi}_t$, $\boldsymbol{\Omega}_t$ and $\boldsymbol{\Psi}_t^x$), but abstracts from government and net export demand, $\boldsymbol{\Psi}_t^g = \boldsymbol{\Psi}_t^{nx} = \mathbf{0}$.

There are five takeaways from Figure 2:

- Consumption shares (red), exceed value-added shares in GDP (black), for both Nondurables and Services while the opposite is true in Construction, Durable Goods, and IPP. Simply put, Nondurable Goods and Services mainly provide final consumption in the U.S. economy while Construction, Durable Goods, and IPP mainly provide inputs for production.

- There is a significant gap between consumption shares and value-added shares for all sectors. This highlights the role of intermediate inputs and capital in determining the overall aggregate importance of different sectors. Adding in production linkages (the series plotted in blue) lessens, but does not eliminate the gap: residual demand from government and net exports remains important for Durables, Nondurables, and Services.

- The rising value-added share of Services, and the corresponding decline in Nondurable Goods, primarily reflect a shift in consumption patterns. Between 1947 and 2018, the Services share of consumption expenditures steadily increased by roughly 22 percentage points, and its share of GDP by 15 percentage points. In contrast, Nondurable Goods as a share of consumption expenditures decreased by 21 percentage points and its share of GDP fell 15 percentage points.

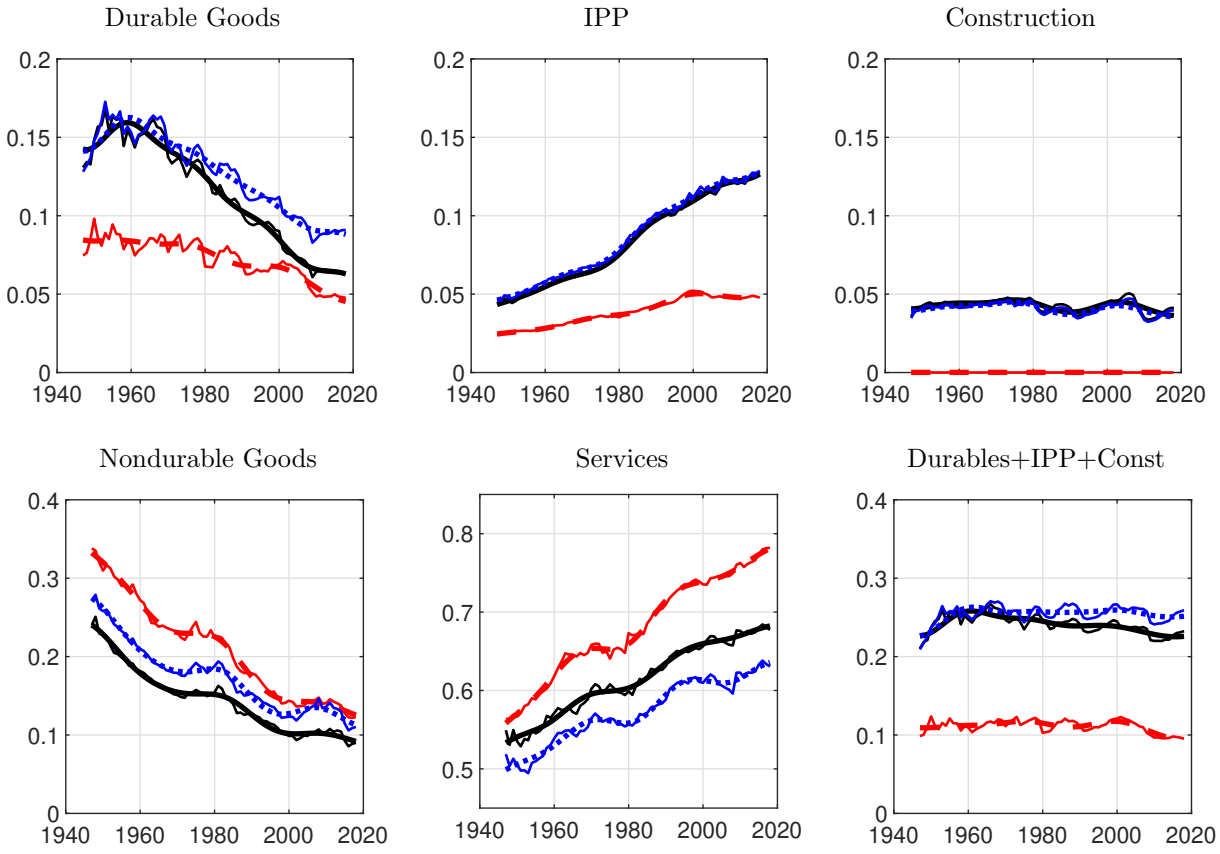
- The share of IPP in GDP has increased steadily during the post-WWII period, and is roughly matched by a decrease in the share of Durables. These trends largely reflect compositional shifts in production rather than consumption and suggest that IPP has essentially displaced Durable Goods as a key source of inputs for production. IPP as a share of GDP today stands where Durable Goods started in 1947.

- The share of Construction in GDP has remained roughly constant.

While goods versus services is a useful taxonomy of sectors, it is also helpful to consider sectors that together contribute mainly to inputs for production (Construction, Durable Goods, IPP) versus those that contribute mainly to final consumption (Nondurable Goods and Services). Viewed in

Figure 2: Trend Decomposition of Sectoral Shares

Consumption Only (red dashed), with Production Linkages (blue dots), Value Added Shares (solid black)



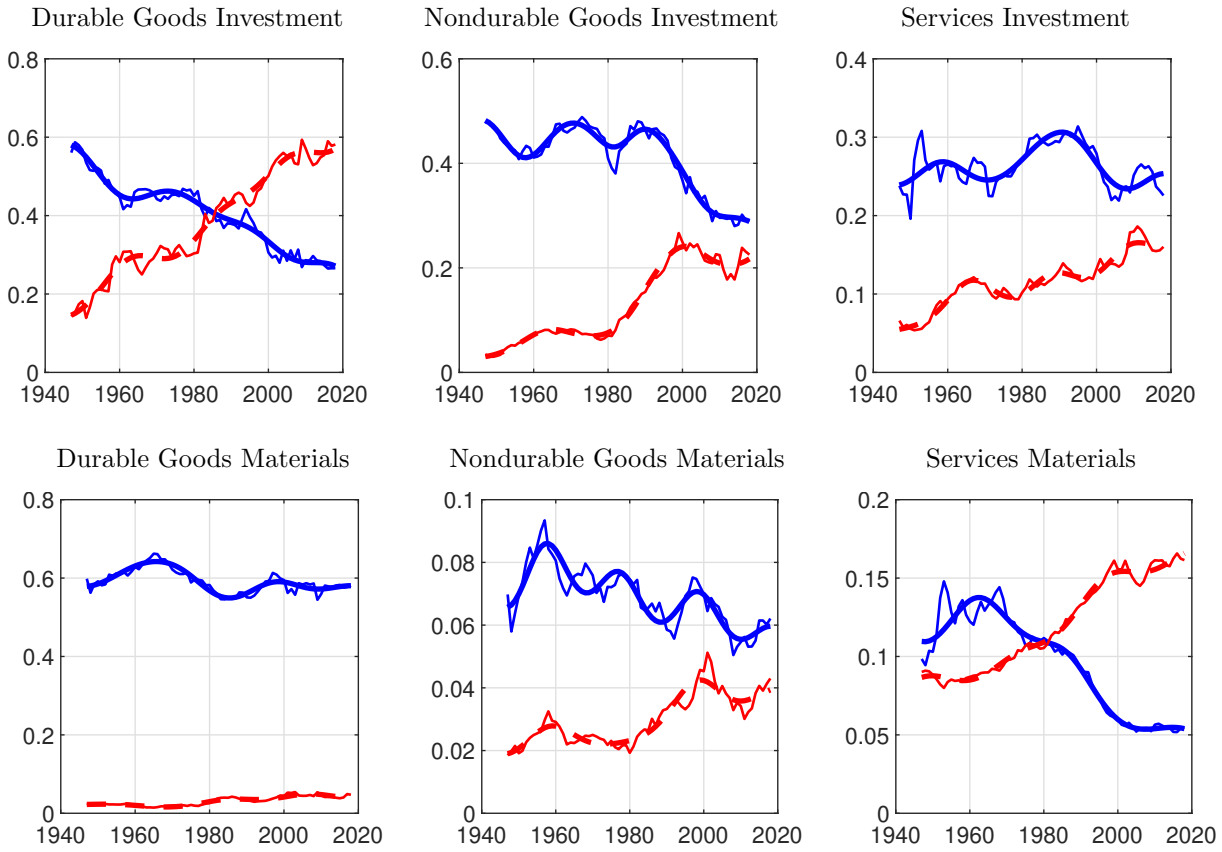
Notes: The red lines describe the evolution of consumption expenditure shares, $\boldsymbol{\theta}_t$. The blue lines add the contributions from production linkages and investment rates to value-added shares in GDP, $\boldsymbol{\Gamma}_t$, $\boldsymbol{\Phi}_t$, $\boldsymbol{\Omega}_t$, and $\boldsymbol{\Psi}_t^x$. Sectoral value-added shares in GDP are shown in black. Thin lines are the raw data and thick lines are low-frequency trends.

this light, the bottom right panel of Figure 2 shows that the value-added share of the combined input-producing sectors vs. those of consumption-producing sectors, including their constituent subshares in the compositional identity (3), has been remarkably constant since 1947. And, since the consumption expenditure shares of Durable Goods and IPP are relatively constant, their change in value-added shares followed from how these sectors contribute to production.

Figure 3 then illustrates some of these secular shifts. The first row of Figure 3 shows IPP and Durable investment expenditure shares for use in the production of Durables, Nondurables, and Services (that is, it shows the corresponding elements of $\boldsymbol{\Omega}_t$, or more precisely the values of $\omega_{ij,t}$ where i indexes IPP and Durables and j indexes Durables, Nondurables, and Services). The top left-hand panel of Figure 3, for example, indicates that in the late 1940s, about 60 percent of investment expenditures by Durable Goods were spent on its own sector, while around 15 percent was spent on IPP. By 2018, investment expenditure shares in Durable Goods had switched. The sector now directed 60 percent of its total investment expenditures toward IPP, and around 25 percent of its

Figure 3: Investment and Materials Spending Shares

Durable Goods (blue) vs. IPP (red), Selected Sectors



Notes: In the top row, the solid blue and dashed red lines describe the evolution of investment spending on Durable Goods and IPP respectively as a share of total investment spending for the different sectors shown. The bottom row depicts the analogous shares for spending on materials. Thin and thick lines depict the raw data and low frequency trends respectively.

investment spending on its own goods. Similar shifts in investment spending are also apparent in Nondurable Goods and Services, though less pronounced. For example, in Transportation, which is a part of our Services definition, motor vehicles continue to play a dominant role as an input even if software is increasingly important to the provision of those services. The bottom row of Figure 3 tells a similar story for intermediate inputs (in this case the corresponding elements of Φ_t are $\phi_{ij,t}$, where i indexes Durable Goods and IPP, and j indexes Durables, Nondurables, and Services), with some exceptions. For example, Durable Goods materials spending shares (the bottom left panel) have remained roughly constant throughout the postwar.

Figure 3 shows a steady overall shift in input use from Durable Goods toward IPP in U.S. production. There is, however, one other striking aspect to this compositional shift: input-cost shares, summed across all inputs, remain approximately constant for each panel in Figure 3. In other words, the switch in input spending from Durables to IPP (as a share of total spending) is virtually one-for-one across all sectors throughout the entire sample period.

Figure 4: International Trade and Government ($\psi^{nx}_{j,t} + \psi^{g}_{j,t}$)

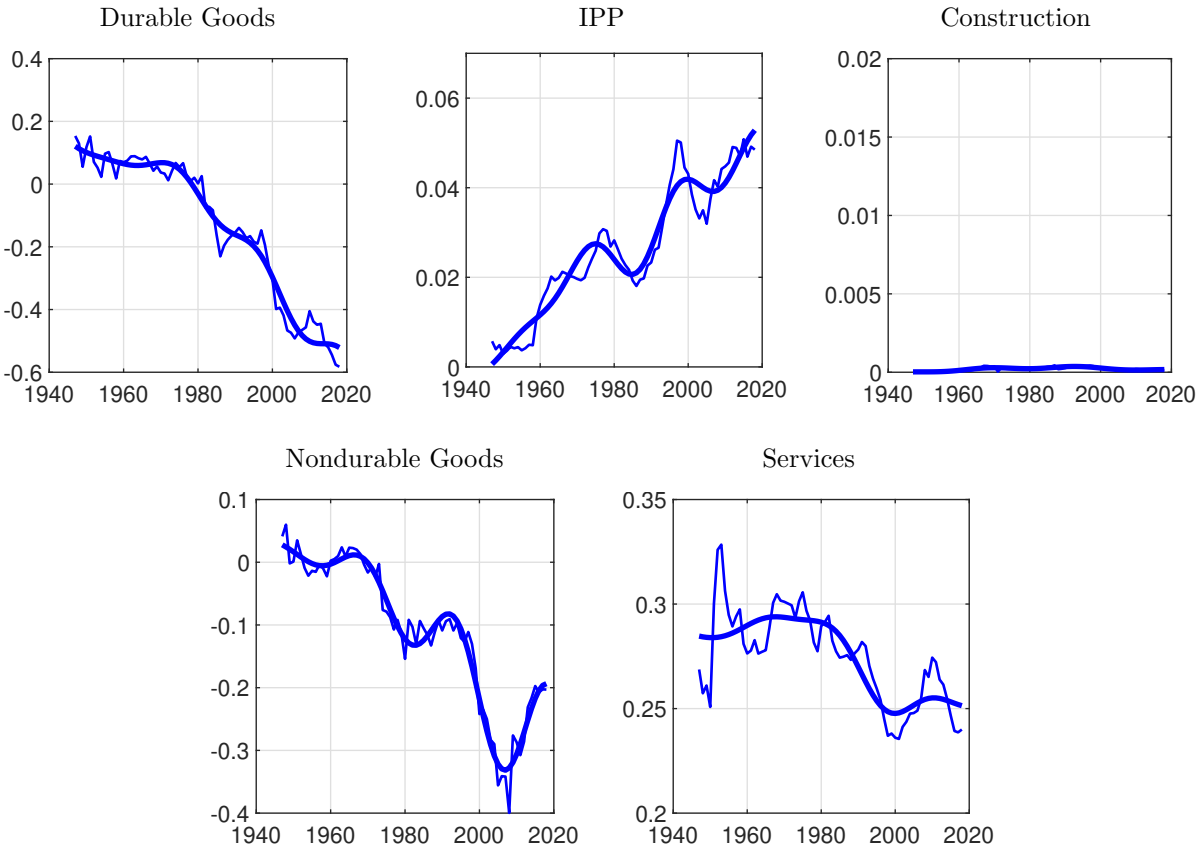


Figure 4 plots the sum, $\psi^{nx}_{j,t} + \psi^{g}_{j,t}$, for each sector and thus shows the compositional role of international trade and government. In all sectors except Services, the government plays no role ($\psi^{g}_{j,t} \approx 0$ for these sectors), so the shares capture changes in international trade. The U.S. was mostly a net exporter of Durable Goods to the rest of the world in the 1950s, 1960s and 1970s but gradually became a net importer.

3.2 Reduced-Form Trend Compositional Forecasting

Having described historical trends in the sectoral composition of the U.S. economy, we now introduce methods for forecasting future trend compositional shifts. We do this in two ways: in this section we use reduced-form statistical methods and in Section 6 we use a structural model described in Sections 4 and 5. We begin with a discussion of methods and then turn to the forecasts.

3.2.1 Forecasting Long-run Trends

We construct long-run forecasts using a modification of methods proposed in Müller and Watson (2016). This subsection includes an overview of the methods; Appendix B and the online Technical Appendix provide a detailed description.

To begin, consider the general problem of predicting the value of the long-run trend in an out-

sample period. Recall that in Section 2.3, we defined the long-run trend, $\hat{x}_t$, as the fitted value from a regression of a time series, x_t, onto a constant, possibly a time trend, and a set of periodic functions, where the regression is estimated using a sample from $t = 1, ..., T$. Thus, suppose that a researcher has data from an 'in-sample' period, say $t = 1, ..., T_{IS}$, and interest focuses on forecasting the trend that would be computed from a longer 'full-sample' period, say $t = 1, ..., T_{FS}$ where $T_{FS} \gg T_{IS}$. Given the construction of $\hat{x}_t$, this requires predicting the value of the sample regression coefficients that would be obtained from the full-sample ($t = 1, ..., T_{FS}$) regression. It turns out that under fairly general assumptions about the stochastic process describing $\{x_t\}$, the full-sample and in-sample regression coefficients are (approximately) jointly normally distributed (see Müller and Watson, 2016). This makes it relatively straightforward to predict the full-sample regression coefficients using the in-sample coefficients.

We operationalize these ideas as follows. First, set $T_{FS} = 2 \times T_{IS}$, so the full-sample period is twice the in-sample period. Second, recall that the parameter, q, governs the periodicities of interest, with the smallest period given by $2T/q$; thus, we set the full-sample value of q at twice its in-sample value, $q_{FS} = 2 \times q_{IS}$, so that the regressions capture the same periodicities. We assume that the process describing x_t is $\Delta x_t = \mu + u_t$ where u_t is an $I(0)$ process; this implies that x_t is $I(1)$ and includes a linear trend unless $\mu = 0$. We estimate the long-run variance of u using a standard HAC ('Newey-West') estimator and ignore sampling uncertainty in the estimated long-run variance when computing the long-run estimates. The method for forecasting a vector-valued long-run trend is a straightforward extension of the scalar problem. Additional details are given in Appendix B and the online Technical Appendix.

In this paper we forecast long-run compositional change in three ways. The first uses the in-sample values of GDP shares shown in Figure 1 to forecast the future trend values. These forecasts use the method described above together with the logit-transformation described in Section 2.3 to impose the support constraints, $\hat{s}_{i,t} \in [0, 1]$ and $\sum_i \hat{s}_{i,t} = 1$, on the forecasts. We refer to these as 'direct forecasts'. The second forecasts begin by forecasting the ingredients of the compositional accounting identity (3), that is, the elements of $(\boldsymbol{\theta}_t, \boldsymbol{\gamma}_t, \boldsymbol{\psi}_t^x, vec(\boldsymbol{\Phi}_t), vec(\boldsymbol{\Omega}_t), \boldsymbol{\psi}_t^g, \boldsymbol{\psi}_t^{nx})$. These are then combined using equation (3) to construct the GDP share forecasts. We refer to these as 'bottom-up forecasts'.[7] Finally, the third forecasts are constructed using the structural model presented in Section 4. We discuss these 'structural-model forecasts' in Section 6.

3.2.2 Long-run Forecasts of Value-Added Shares

Table 2 summarizes the findings for direct forecasts of 20-year ahead trend value-added shares. The resulting forecasts show the continuation of historical trends over the next 20 years. Services and IPP continue to increase in relative importance to the U.S. economy while Durables and Nondurables continue to decline, and Construction remains largely unchanged. However, compared to historical average 20-year changes, the table indicates a slowdown in the evolution of these shares. Services' share of GDP, for example, has historically increased by approximately 4 percentage points on

[7] The Online Apprendix discusses the forecasts of $(\boldsymbol{\theta}_t, \boldsymbol{\gamma}_t, \boldsymbol{\psi}_t^x, vec(\boldsymbol{\Phi}_t), vec(\boldsymbol{\Omega}_t), \boldsymbol{\psi}_t^g, \boldsymbol{\psi}_t^{nx})$.

Table 2: Forecasts of Trend Sectoral Value Added Shares of GDP
2018 to 2038 (Percentage Points)

Sector	Durable Goods	IPP	Construction	Nondurable Goods	Services
Levels					
Actual 1947	14.2	4.3	4.1	24	53.3
Actual 2018	6.3	12.7	3.6	9.2	68.2
Forecast for 2038	4.7	16.1	3.3	6.6	69.2
	[3.7,5.9]	[14.9,17.3]	[2.5,4.5]	[5.5,7.9]	[66.9,70.7]
Changes					
Actual 1947-2018	-8.0	8.3	-0.4	-14.9	14.9
20 Year Hist Avg	-2.2	2.3	-0.1	-4.1	4.1
Forecast 2018-2038	-1.5	3.5	-0.3	-2.6	0.9
	[-2.4,-0.3]	[2.3,4.5]	[-1.1,0.7]	[-3.6,-1.4]	[-1.1,2.5]

Notes: The forecasts are the means of the associated predictive distributions. Values in brackets are equal-tailed 67% prediction intervals.

average every 20-year period, but its mean trend forecast over the next 20 years indicates less than a 1 percentage point increase. Similarly, the pace of the trend decline in the relative importance of Durable Goods and Nondurable Goods is forecasted to abate over the next 20 years.

In addition to the expected slowdown in the pace of future structural change, a key difference between expected changes over the next 20 years and historical 20-year changes in Table 2 is in the pattern of compositional change. Historically, changes in the shares of Services and Nondurable Goods offset each other, at 4.1 and -4.1 percentage points respectively on average. Similarly, changes in the shares of IPP and Durable Goods also offset each other historically, at 2.3 and -2.2 percentage points respectively. This is not so for the reduced-form forecasts. While the relative size of Services is expected to continue to rise over the next 20 years, this increase falls considerably short in absolute magnitude of the expected decline in Nondurable Goods over the same period. Similarly, since share changes must sum to zero across sectors, the pace at which IPP's relative size is expected to increase in the next 20 years far exceeds the expected decline in the relative size of Durables.

We note that the form of the logit transformation, L, plays an important role in these forecasts: since shares cannot fall below 0 or exceed 1, the fall in the expected pace of low-frequency changes in sectoral shares is governed by the way in which the logit transformation, L, imposes these constraints. Thus, with an already small share of GDP, the decline in the relative size of Durable Goods is expected to slow going forward. In contrast, with a share of GDP around 13 percent, the relative size of IPP, purely from the standpoint of respecting its upper bound constraint, has considerable space to increase further. In Section 6, the logit transformation is replaced by a structural model where the share adding up constraints arises endogenously.

Results for the bottom-up forecasts are summarized in Table 3. The overall sectoral changes

Table 3: Decomposing Forecasts of Changes in Sectoral Value Added Shares
2018 to 2038 (Percentage Points)

Sectors	Durables	IPP	Construction	Nondurables	Services
Forecast Changes					
Direct	-1.5	3.5	-0.3	-2.6	0.9
Bottom Up	-1.6	3.5	-0.4	-2.6	1.1
By Component:					
$\boldsymbol{\theta}_t$	-0.4	0.3	0.0	-1.3	1.4
$\boldsymbol{\gamma}_t$	0.0	0.3	0.2	0.1	-0.7
$\boldsymbol{\Phi}_t$	-0.3	1.0	-0.4	-0.1	-0.1
$\boldsymbol{\Omega}_t$	-0.5	1.6	-0.1	-1.2	0.2
$\boldsymbol{\psi}_t^X$	0.0	0.2	-0.2	0.1	-0.1
$\boldsymbol{\psi}_t^{NX}$	-0.5	0.1	0.0	-0.3	0.6
$\boldsymbol{\psi}_t^G$	0.0	0.1	0.0	0.1	-0.2

Notes: Adding up is subject to rounding and non-linearities in compositional identities.

predicted by the bottom-up approach are shown in panel (a) and are quite similar to the direct forecasts. Panel (b) uses the compositional accounting identity (3) to decompose the predicted overall changes. It shows that, as with historical trends, the forecasts of long-term compositional changes in Services and Nondurable Goods are dominated by persistent shifts in the shares of expenditure on consumption, $\boldsymbol{\theta}_t$. Similarly, the continued increase in the importance of IPP in GDP, and corresponding decline in that of Durable Goods, is primarily driven by continued shifts in the mix of input use, as reflected in $\boldsymbol{\Omega}_t$ and $\boldsymbol{\Phi}_t$, though net exports still play a notable role in shaping the future of Durable Goods and Services.

Summing up the results from this section, the compositional identity (3) helped us break down some of the key aspects of the low-frequency rise and fall in the relative importance of different U.S. production sectors; this was a compositional accounting exercise. This reduced-form compositional framework also gave us a sense of how these historical trends might evolve over the next 20 years; this was a compositional forecasting exercise. That said, both of these reduced-form exercises have inherent limitations. Compositional accounting, while relating overall sectoral value added shares of GDP to underlying constituent subshares (e.g., consumption expenditure shares or input cost shares), cannot tell us how these constituent shares reflect more fundamental forces such as sectoral productivity changes or changes in preferences. At the same time, reduced-form compositional forecasts reflect in part an ad-hoc logit transformation rather than forecasts of the same underlying fundamentals in a framework where shares endogenously satisfy their support constraints. Thus, we now turn to addressing these limitations.

4 Modeling Structural Change

The previous section highlighted key changes in consumption patterns and the mix of inputs used in production as largely accounting for observed secular compositional shifts in U.S. GDP since WWII. For the purpose of modeling structural change, three observations stand out.

First, the consumption expenditure share of Nondurables steadily fell over time, but in a way that was essentially offset by a rise in the consumption share of Services, leaving their combined importance in total consumption expenditures largely unchanged. In addition, Figure 2 showed that the change in the share of Nondurable consumption (red line) accounted for the major part of the change in its value added share of GDP (black line), and similarly for Services.

Second, production in the U.S., regardless of the sector, gradually moved from employing Durable Goods to employing IPP as inputs. Recall in Figure 2 that in Construction, Durable Goods, and IPP, changes in consumption shares were relatively small, while the value added share of Durable Goods steadily declined, and that of IPP steadily increased. Figure 3 further showed a clear switch in sectors' input spending from Durable Goods to IPP. Remarkably, trend shifts in input cost shares related to Durable Goods and IPP essentially offset each other in all sectors and throughout the entire sample.

Finally, an important implication of these changes is that while some shares have steadily shifted over the postwar period − for example, input cost shares related to IPP and Durable Goods or consumption expenditure shares associated with Nondurable Goods and Services − other shares have remained remarkably constant throughout. This is the case, for instance, for the combined shares of Construction, Durable Goods, and IPP in GDP, their combined consumption expenditure shares (bottom right panel of Figure 2), or the IPP and Durable Goods materials cost shares in the Durable Goods sector (bottom left panel of Figure 3).

To address these observations, we present a model that generically partitions the set of all sectors, $\mathcal{N}$, into two subsets. We denote by $\widetilde{\mathcal{N}}$ the set of sectors whose goods are assembled using constant-elasticity-of-substitution (CES) aggregators, and whose shares will vary with equilibrium prices and quantities over time. We denote by $\bar{\mathcal{N}}$ the set of sectors whose goods are assembled using unit-elastic aggregators, and whose shares will be constant and unrelated to equilibrium allocations. These two subsets cover all sectors, $\mathcal{N} = \bar{\mathcal{N}} \cup \widetilde{\mathcal{N}}$, and are disjoint, $\bar{\mathcal{N}} \cap \widetilde{\mathcal{N}} = \emptyset$.

Importantly, the subsets $\widetilde{\mathcal{N}}$ or $\bar{\mathcal{N}}$ will not always include the same sectors depending on which aspect of the environment, production or preferences, is being addressed. For example, we let $\widetilde{\mathcal{N}}^x$ denote the set of sectors whose goods are assembled into investment using a CES technology. Thus, as suggested by Figure 3, $\widetilde{\mathcal{N}}^x$ will include Durable Goods and IPP. Similarly, we denote by $\widetilde{\mathcal{N}}^c$ the set of sectors whose goods are bundled into consumption using a (generalized) CES aggregator. In this case, the consumption shares in Figure 2 suggest defining $\widetilde{\mathcal{N}}^c$ as including Nondurable Goods and Services.

Our interest lies in secular changes across several decades, and we therefore abstract from transition dynamics in interpreting trends. We interpret these trends as a sequence of steady states, rather than as a balanced growth path, for two reasons. First, the observed secular compositional

changes in value added, consumption, or input cost shares appear at face value inconsistent with balanced growth. Even though Ngai and Pissarides (2007), and more recently Gaggl et al. (2023), show that models with an aggregate balanced growth path can be consistent with changing sectoral shares, we do not follow this approach because it requires particular assumptions on preferences and production, and restrictions on growth rates across sectors.[8] Second, in Foerster et al. (2022) we argue that the underlying trends for sectoral TFP growth rates, and resulting real GDP growth, have not remained constant but generally declined between 1950 and 2018.

We then use our structural analysis to address the underlying sources of compositional change, something that could not be addressed in the reduced form compositional accounting framework described in the previous section. We begin by quantifying how different forces, both on the supply side and demand side, have contributed to structural change in general equilibrium. Exogenous supply-side changes arise in the guise of changes in total factor productivity (TFP). They also arise through input-biased technical change (IBTC): the efficiency with which inputs translate into output in different sectors. Exogenous changes on the demand side are associated with changes in preferences and net exports.

We then construct long-run forecasts of compositional change that arise in the structural model. Two key differences distinguish structural compositional forecasting from the reduced-form approach adopted earlier. First, although the compositional identity (3) must also hold in the structural model, the structural forecasts are driven by forecasts of fundamental forces, e.g., forecasts of sectoral trends in TFP or forecasts of exogenous preference trends. Second, the somewhat ad hoc logit transformation that ensured that the reduced-form compositional forecasts satisfied their support constraints is no longer needed. Instead, the layout of the structural model, including the choices of technology and preferences, enforce these constraints endogenously.

4.1 Technology

At each date t, gross output in each sector, j, $y_{j,t}$, is assembled using a value-added bundle, $v_{j,t}$, and a materials bundle, $m_{j,t}$, using a Cobb-Douglas technology,

$$y_{j,t} = \mathcal{Y}\left(v_{j,t}, m_{j,t}\right) = \left(\frac{v_{j,t}}{\gamma_j}\right)^{\gamma_j} \left(\frac{m_{j,t}}{1-\gamma_j}\right)^{1-\gamma_j}, \ \gamma_j \in (0,1), \ \forall j \in \mathcal{N}, \tag{4}$$

with implied price index,

$$p^y_{j,t} = \mathcal{P}^Y_j\left(p^v_{j,t}, p^m_{j,t}\right) = \left(p^v_{j,t}\right)^{\gamma_j} \left(p^m_{j,t}\right)^{1-\gamma_j},$$

where $p^v_{j,t}$ and $p^m_{j,t}$ are the price indices for the value-added and materials bundles respectively. The gross-output technology, $\mathcal{Y}\left(v_{j,t}, m_{j,t}\right)$, implies constant value-added and materials shares in every sector, $\frac{p^v_{j,t} v_{j,t}}{p^y_{j,t} y_{j,t}} = \frac{V_{j,t}}{Y_{j,t}} = \gamma_j$ and $\frac{p^m_{j,t} m_{j,t}}{p^y_{j,t} y_{j,t}} = \frac{M_{j,t}}{Y_{j,t}} = 1 - \gamma_j$, respectively $\forall j \in \mathcal{N}$.

The materials bundle used by sector j, $m_{j,t}$, is also produced with a unit-elastic technology that

[8] Essentially, the aggregate balanced growth rate is defined with respect to aggregates defined in terms of capital goods, and its existence requires a common capital good for all sectors and indirect preferences consistent with a constant interest rate when expenditures grow at a constant rate (Herrendorf et al., 2021).

uses two materials sub-bundles, $\widetilde{m}_{j,t}$ and $\bar{m}_{j,t}$, according to

$$m_{j,t} = \mathcal{M}(\widetilde{m}_{j,t}, \bar{m}_{j,t}) = \left(\frac{\widetilde{m}_{j,t}}{\rho_j^m}\right)^{\rho_j^m} \left(\frac{\bar{m}_{j,t}}{1-\rho_j^m}\right)^{1-\rho_j^m}, \quad \rho_j^m \in (0,1),\ \forall j \in \mathcal{N}. \tag{5}$$

The materials sub-bundles, $\widetilde{m}_{j,t}$ and $\bar{m}_{j,t}$, however, are assembled using different technologies depending on the sectors from which their inputs are sourced,

$$\widetilde{m}_{j,t} = \left[\sum_{i \in \widetilde{\mathcal{N}}_j^m} z_{ij,t}^m m_{ij,t}^{\frac{\epsilon_j^m - 1}{\epsilon_j^m}}\right]^{\frac{\epsilon_j^m}{\epsilon_j^m - 1}}, \quad \sum_{i \in \widetilde{\mathcal{N}}_j^m} z_{ij,t}^m = 1,\ \epsilon_j^m \in (0,\infty), \tag{6}$$

$$\bar{m}_{j,t} = \prod_{i \in \bar{\mathcal{N}}_j^m} \left(\frac{m_{ij,t}}{\zeta_{ij}^m}\right)^{\zeta_{ij}^m}, \quad \sum_{i \in \bar{\mathcal{N}}_j^m} \zeta_{ij}^m = 1, \tag{7}$$

where $m_{ij,t}$ denotes sector i goods used as intermediate inputs by sector j. Concretely, given the findings in Section 3, $\widetilde{\mathcal{N}}_j^m$ includes, for every sector j, Durable Goods and IPP while $\bar{\mathcal{N}}_j^m$ includes Construction, Nondurables, and Services.[9] The cost-minimization problem faced by the representative firm producing $m_{j,t}$ implies a materials price index, $p_{j,t}^m$, given by

$$p_{j,t}^m = \mathcal{P}_j^M(p_t^y) = \left[\sum_{i \in \widetilde{\mathcal{N}}_j^m} (z_{ij,t}^m)^{\epsilon_j^m} \left(p_{i,t}^y\right)^{1-\epsilon_j^m}\right]^{\frac{\rho_j^m}{1-\epsilon_j^m}} \prod_{i \in \bar{\mathcal{N}}_j^m} (p_{i,t}^y)^{\zeta_{ij}^m (1-\rho_j^m)} \quad \forall j \in \mathcal{N},$$

where $p_{i,t}^y$ is the (gross output) price of sector i goods and $p_t^y = (p_{1,t}^y, ..., p_{n,t}^y)'$.

The functional form for the production of the sub-bundle, $\widetilde{m}_{j,t}$, means that sector j's expenditure shares on materials from sector i goods, $\frac{p_{i,t}^y m_{ij,t}}{p_{j,t}^m m_{j,t}} = \frac{M_{ij,t}}{M_{j,t}}$, $i \in \widetilde{\mathcal{N}}_j^m$, can shift as prices change over time. Conversely, the production of the materials sub-bundle, $\bar{m}_{j,t}$, means that the material spending shares, $\frac{M_{ij,t}}{M_{j,t}}$, $i \in \bar{N}_j^m$, will be time invariant. In practical terms, therefore, this production structure allows any sector, j, to vary its relative spending shares on materials between Durable Goods and IPP as shown in Figure 3. However, the specification of the materials aggregate, $\mathcal{M}(.)$, implies that the sum of these two input cost shares will be constant and equal to ρ_j^m. Finally, $z_{ij,t}^m$ captures the relative efficiency of inputs from sector $i \in \widetilde{\mathcal{N}}_j^m$, specifically Durable Goods relative to IPP, when used in production by sector j: input-biased technical change (IBTC).

The value added from capital, $k_{j,t}$, and labor, $\ell_{j,t}$, is produced using a unit-elasticity technology in all sectors,

$$v_{j,t} = \mathcal{V}_j(k_{j,t}, \ell_{j,t}) = z_{j,t} \left(\frac{k_{j,t}}{\alpha_j}\right)^{\alpha_j} \left(\frac{\ell_{j,t}}{1-\alpha_j}\right)^{1-\alpha_j}, \ \alpha_j \in (0,1),\ j \in \mathcal{N}, \tag{8}$$

[9]The model notation implies the CES bundles in materials, $\bar{\mathcal{N}}_j^m$, and in investment $\bar{\mathcal{N}}_j^x$, can contain different industries for each j. However, as shown in Section 3, the industries in these bundles are the same in our application.

where $z_{j,t}$ captures sector-specific total factor productivity (TFP). The associated rental and wage rates are $u_{j,t} = p^v_{j,t} z_{j,t} \left(\frac{k_{j,t}}{\alpha_j}\right)^{\alpha_j - 1} \left(\frac{\ell_{j,t}}{1-\alpha_j}\right)^{1-\alpha_j}$ and $w_{j,t} = p^v_{j,t} z_{j,t} \left(\frac{k_{j,t}}{\alpha_j}\right)^{\alpha_j} \left(\frac{\ell_{j,t}}{1-\alpha_j}\right)^{-\alpha_j}$, respectively.

Capital in sector j depreciates at rate $\delta_j \in (0,1)$ and accumulates according to $k_{j,t+1} = x_{j,t} + (1-\delta_j)k_{j,t}$, where $x_{j,t}$ denotes sector j investment. As with materials, sector j investment is produced by combining two sub-bundles, $\widetilde{x}_{j,t}$ and $\bar{x}_{j,t}$, using a unit-elastic technology,

$$x_{j,t} = \mathcal{X}(\widetilde{x}_{j,t}, \bar{x}_{j,t}) = \left(\frac{\widetilde{x}_{j,t}}{\rho^x_j}\right)^{\rho^x_j} \left(\frac{\bar{x}_{j,t}}{1-\rho^x_j}\right)^{1-\rho^x_j}, \quad \rho^x_j \in (0,1),\ j \in \mathcal{N}. \tag{9}$$

The sub-bundles, $\widetilde{x}_{j,t}$ and $\bar{x}_{j,t}$, are in turn assembled using different technologies depending on the sectors from which their inputs are acquired,

$$\widetilde{x}_{j,t} = \left[\sum_{i \in \widetilde{\mathcal{N}}^x_j} z^{x\ \frac{\epsilon^x_j - 1}{\epsilon^x_j}}_{ij,t} x^{\frac{\epsilon^x_j - 1}{\epsilon^x_j}}_{ij,t}\right]^{\frac{\epsilon^x_j}{\epsilon^x_j - 1}}, \ \sum_{i \in \widetilde{\mathcal{N}}^x_j} z^x_{ij,t} = 1,\ \epsilon^x_j \in (0,\infty), \tag{10}$$

$$\bar{x}_{j,t} = \prod_{i \in \bar{\mathcal{N}}^x_j} \left(\frac{x_{ij,t}}{\zeta^x_{ij}}\right)^{\zeta^x_{ij}}, \ \sum_{i \in \bar{\mathcal{N}}^x_j} \zeta^x_{ij} = 1, \tag{11}$$

where $x_{ij,t}$ denotes sector i goods used as investment by sector j. The associated cost-minimization problem yields a price index for sector j investment, $p^x_{j,t}$, given by

$$p^x_{j,t} = \mathcal{P}^X_j(p^y_t) = \left[\sum_{i \in \widetilde{\mathcal{N}}^x_j} (z^x_{ij,t})^{\epsilon^x_j} \left(p^y_{i,t}\right)^{1-\epsilon^x_j}\right]^{\frac{\rho^x_j}{1-\epsilon^x_j}} \prod_{i \in \bar{\mathcal{N}}^x_j} (p^y_{i,t})^{\zeta^x_{ij}(1-\rho^x_j)}.$$

Similarly to materials, these functional forms allow sector j's shares of investment spending on sector i goods, $\frac{p^y_{i,t} x_{ij,t}}{p^x_{j,t} x_{j,t}} = \frac{X_{ij,t}}{X_{j,t}}$, $i \in \widetilde{\mathcal{N}}^x_j$, to change with prices over time while allowing the investment spending shares, $\frac{X_{ij,t}}{X_{j,t}}$, $i \in \bar{\mathcal{N}}^x_j$, to be time-invariant. Moreover, guided by the findings in Section 3, the set $\widetilde{\mathcal{N}}^x_j$ includes Durable Goods and IPP so that, consistent with Figure 3, investment spending shares of these two sectors may shift over time. Finally, as with the technology for materials production, $z^x_{ij,t}$ denotes the efficiency of new capital sourced from Durable Goods relative to that of IPP in the production of sector j output.

4.2 Preferences

Similar to the production side, we distinguish between two sets of consumption goods, $j \in \widetilde{\mathcal{N}}^c$ and $j \in \bar{\mathcal{N}}^c$. Goods $j \in \widetilde{\mathcal{N}}^c$ are bundled using a non-homothetic aggregator such that their shares in total consumption expenditures will vary with relative prices and the level of total expenditures. Goods $j \in \bar{\mathcal{N}}^c$ are bundled using a unit-elasticity aggregator and will have constant shares. This breakdown is motivated by the behavior of consumption shares in Figure 2 (the red lines) where

$\widetilde{\mathcal{N}}^c$ includes Nondurables and Services while $\bar{\mathcal{N}}^c$ comprises Construction, Durable Goods, and IPP. Total household expenditures on consumption, e_t, reflect expenditures on both types of goods with expenditure on goods $j \in \widetilde{\mathcal{N}}^c$ and $j \in \bar{\mathcal{N}}^c$ denoted by $\widetilde{e}_t$ and $\bar{e}_t$ respectively, so that $e_t = \widetilde{e}_t + \bar{e}_t$.

Following Comin, Mestieri, and Lashkari (2021), the utility aggregator, $\widetilde{c}_t = \mathcal{C}(c_{1,t}, ..., c_{n,t})$, is defined implicitly by

$$\sum_{j \in \widetilde{\mathcal{N}}^c} \Theta_{j,t}^{1/\sigma} \left[c_{j,t} / (\widetilde{c}_t)^{\epsilon_j^c} \right]^{(\sigma-1)/\sigma} = 1, \tag{12}$$

where $\Theta_{j,t}$ captures exogenous shifts in preferences for particular goods. The preferences described in equation (12) are of interest here because they allow changes in relative prices and income to affect consumption expenditure shares. Comin et al. (2021) show that the parameters, $\Theta_{j,t}$, and, ϵ_j^c, are not identified from observations on prices and expenditures absent a normalization with respect to some arbitrary base commodity, b. Thus, we set $\Theta_{b,t} = \epsilon_b = 1 \; \forall t$ for some commodity $b \in \widetilde{\mathcal{N}}^c$ and treat equation (12) as the normalized version of preferences. All other preference parameters satisfy $\Theta_{j,t}, \epsilon_j \geq 0$, and $\sigma \geq 0$. The associated expenditure minimization problem yields expenditures,

$$\widetilde{e}_t = \mathcal{E}(p_t^y, \widetilde{c}_t) = \left[\sum_{j \in \widetilde{\mathcal{N}}^c} \Theta_{j,t} (p_{j,t}^y)^{1-\sigma} (\widetilde{c}_t)^{\epsilon_j^c (1-\sigma)} \right]^{\frac{1}{1-\sigma}}. \tag{13}$$

In sectors $j \in \bar{\mathcal{N}}^c$, goods are bundled into a homothetic index according to the unit-elastic aggregator,

$$\bar{c}_t = \prod_{j \in \bar{\mathcal{N}}^c} \left(\frac{c_{j,t}}{\zeta_j^c} \right)^{\zeta_j^c}, \quad \sum_{j \in \bar{\mathcal{N}}^c} \zeta_j^c = 1, \tag{14}$$

with corresponding consumption expenditures, $\bar{e}_t = p_t^{\bar{c}} \bar{c}_t$, where $p_t^{\bar{c}} = \prod_{j \in \bar{\mathcal{N}}^c} (p_{j,t}^y)^{\zeta_j^c}$ is the price index for the homothetic bundle. The representative household then solves the utility maximization problem,

$$\max \sum_{t=0}^{\infty} \beta^t \left[(\bar{c}_t)^{\rho_t^c} (\widetilde{c}_t)^{1-\rho_t^c} - \sum_{j \in \mathcal{N}} \frac{\varphi_{j,t} \ell_{j,t}^{1+\gamma_\ell}}{1+\gamma_\ell} \right] \tag{15}$$

subject to

$$p_t^{\bar{c}} \bar{c}_t + \mathcal{E}(p_t^y, \widetilde{c}_t) + \sum_{j \in \mathcal{N}} p_{j,t}^x \left[k_{j,t+1} - (1-\delta_j) k_{j,t} \right] = \sum_{j \in \mathcal{N}} w_{j,t} \ell_{j,t} + \sum_{j \in \mathcal{N}} u_{j,t} k_{j,t}, \tag{16}$$

where $\ell_{j,t}$ and $\varphi_{j,t}$ denote, respectively, labor input and labor supply shifters in sector j while γ_ℓ is the Frisch elasticity of labor.[10]

[10] We allow for time variation in the relative weight of the (non)homothetic consumption index, ρ_t^c, in equation (15). With ρ^c fixed, the non-homotheticity of preferences implies that, in equilibrium, the consumption share of the homothetic bundle, $\bar{c}_t$, changes over time while Figure 2 indicates that it is largely constant. Thus, we assume some compensating variation in ρ_t^c (see the online Technical Appendix).

4.3 Steady State Equilibrium Prices, Allocations, and Shares

Because the evolution of trend shares discussed in Section 3 is defined in terms of (approximate) 20 year variations, our analysis focuses on sequential steady states driven by trend variations in supply and demand forces. On the supply side, the exogenous processes that drive equilibrium prices and quantities are sectoral TFP, $z_{j,t}$, input-biased technical change (IBTC) in intermediate inputs and investment, $z^m_{ij,t}$ and $z^x_{ij,t}$, respectively, and labor supply shifts, $\varphi_{j,t}$. Demand-side driving processes include exogenous shifts in preferences, $\Theta_{j,t}$, and changes in the ratio of government spending and net exports to value-added, $\psi^{nx}_{j,t}$ and $\psi^{g}_{j,t}$. That is, we do not present an explicit theory of changes in trade but trace out their effects through the economy's production networks. We define the low-frequency trend in these various driving processes as in Section 3. We then interpret secular shifts in prices and allocations as sequential steady states of our model economy driven by these low frequency trends.[11] Thus, while we define the steady state equilibrium (SSE) below without time subscripts, we retain the their use thereafter where appropriate.

Steady State Equilibrium: *For all $i, j \in \mathcal{N}$, given z_j, z^m_{ij}, z^x_{ij}, φ_j, Θ_j, ψ^g_j, and ψ^{nx}_j, a steady state equilibrium (SSE) is defined by a set of sectoral prices, $\{p^y_j, p^v_j, p^m_j, p^x_j, u_j, w_j\}$, and quantities, $\{c_j, x_{ij}, x_j, m_{ij}, m_j, \ell_j, k_j, v_j, y_j\}$, such that, taking prices as given, firms maximize profits, households maximize utility and all markets clear.*

The online Technical Appendix describes in detail the steps involved in solving for steady state prices and quantities at any date t. Ultimately, solving for steady state prices and quantities reduces to solving a set of n residual excess demand equations in n unknowns, namely gross output prices, $p^y_{j,t}$ $j \in \mathcal{N}$. The online Technical Appendix shows that although the non-homothetic component of preferences is not homogeneous of degree one in its real consumption index, $\widetilde{c}$, the corresponding excess demand functions continue to be homogeneous of degree one in prices. In other words, there are only $n-1$ independent residual equations and one needs to choose a numéraire, which we take to be Nondurable Goods.

The environment is set up in such a way that, in equilibrium, most shares including input cost shares of materials and investment, consumption shares, etc. will be independent of low-frequency shifts in prices and quantities. Put another way, an important part of the model remains anchored by standard unit-elastic assumptions and is straightforward to quantify. A direct implication is that consumption shares related to Construction, Durable Goods, and IPP, or input cost shares related to Construction, Nondurable Goods, and Services, will all be time-invariant. However, in other key areas of production or consumption, the composition of inputs or consumption will change as equilibrium prices and quantities respond to secular changes in the driving processes. Thus, we now turn to the model's low-frequency implications for the various shares described in Section 3.

In every sector j, equilibrium input cost shares pertaining to materials purchased from sector

[11] This approximation is motivated by the notion that transition dynamics are most pertinent for business cycle dynamics rather than low frequency trends; see Foerster et al. (2022) for an alternative application.

$i \in \bar{\mathcal{N}}_j^m$ are given by

$$\frac{p_{i,t}^y m_{ij,t}}{p_{j,t}^m m_{j,t}} = \zeta_{ij}^m (1-\rho_j^m), \forall j \text{ and } \; i \in \bar{\mathcal{N}}_j^m. \tag{17}$$

In other words, in every sector, the cost shares associated with materials purchased from Construction, Nondurables, and Services are independent of allocations or prices. Therefore, even if sectoral allocations or prices slowly shift over time, these cost shares will remain constant, though not necessarily the same in each sector j. In the input-output matrix, $\mathbf{\Phi}_t$, they represent time-invariant rows, ϕ_{ij}, associated with $i \in \bar{\mathcal{N}}_j^m \; \forall j$.

For materials purchased from sectors $i \in \widetilde{\mathcal{N}}_j^m$, we have

$$\frac{p_{i,t}^y m_{ij,t}}{p_{j,t}^m m_{j,t}} = (\widetilde{m}_{j,t})^{\frac{1-\epsilon_j^m}{\epsilon_j^m}} (m_{ij,t})^{\frac{\epsilon_j^m - 1}{\epsilon_j^m}} z_{ij,t}^m \rho_j^m, \forall j \text{ and } \; i \in \widetilde{\mathcal{N}}_j^m. \tag{18}$$

Thus, in every sector, j, the cost shares corresponding to materials sourced from Durable Goods and IPP will shift with secular changes in allocations and, importantly, input biased technical change, $z_{ij,t}^m$. Furthermore, in addition to IBTC, in general equilibrium, low-frequency changes in TFP or preferences influence prices and, therefore, the associated quantities, $\widetilde{m}_{j,t}$ and $m_{ij,t}$, used in the production sector j goods. In particular, we can rewrite equation (18) for materials sourced from $i, i' \in \widetilde{\mathcal{N}}_j^m$ as,

$$\log \frac{M_{ij,t}}{M_{i'j,t}} = \log \frac{p_{i,t}^y m_{ij,t} / p_{j,t}^m m_{j,t}}{p_{i',t}^y m_{i'j,t} / p_{j,t}^m m_{j,t}} = \epsilon_j^m \log \frac{z_{ij,t}^m}{z_{i'j,t}^m} + (1-\epsilon_j^m) \log \frac{p_{i,t}^y}{p_{i',t}^y}. \tag{19}$$

Therefore, in all sectors j, changes in relative spending on materials from, say, Durable Goods and IPP, $M_{ij,t}/M_{i'j,t}$, reflect changes in the price of Durable Goods relative to that of IPP, $p_{i,t}^y / p_{i',t}^y$, and the efficiency with which IPP materials translate into output relative to that of Durable Goods, $z_{ij,t}^m / z_{i'j,t}^m$. The response of relative cost shares to changes in relative prices in sector j depends on the extent to which materials composed of Durable Goods and IPP act as substitutes or complements in that sector, $\epsilon_j^m \lessgtr 1$. Figure 3 and equation (19) together suggest a secular increase in the productivity of materials sourced from IPP compared to Durable Goods for most industries over the postwar period.

Finally, it is important to recognize that while we allow some input cost shares to shift with prices and allocations at low frequencies, cost shares for the overall bundles defined by $\bar{\mathcal{N}}_j^m$ and $\widetilde{\mathcal{N}}_j^m$ are independent of these prices and allocations and time invariant, $\sum_{i \in \mathcal{N}} \frac{p_{i,t}^y m_{ij,t}}{p_{j,t}^m m_{j,t}} = \sum_{i \in \widetilde{\mathcal{N}}_j^m} \frac{p_{i,t}^y m_{ij,t}}{p_{j,t}^m m_{j,t}} + \sum_{i \in \bar{\mathcal{N}}_j^m} \frac{p_{i,t}^y m_{ij,t}}{p_{j,t}^m m_{j,t}} = \rho_j^m + (1-\rho_j^m) = 1$. Put another way, we keep departures from unit-elastic assumptions limited as they apply only to individual inputs sourced from within $\widetilde{\mathcal{N}}_j^m$.

Analogously to materials, equilibrium investment spending shares in all sectors j sourced from sectors $i \in \bar{\mathcal{N}}_j^x$ are

$$\frac{p_{i,t}^y x_{ij,t}}{p_{j,t}^x x_{j,t}} = \zeta_{ij}^x (1-\rho_j^x), \forall j \text{ and } \; i \in \bar{\mathcal{N}}_j^x,$$

and have similar properties. Thus, the cost shares associated with investment spending on Construction, Nondurables, and Services are independent of equilibrium prices and allocations but can differ across sectors. They represent the time-invariant rows, $\omega_{ij},\ i \in \bar{\mathcal{N}}_j^x$, of the investment flow matrix, $\boldsymbol{\Omega}_t$.

For new capital purchased from sectors $i, i' \in \widetilde{\mathcal{N}}_j^x$, we have that

$$\log \frac{p_{i,t}^y x_{ij,t}/p_{j,t}^x x_{j,t}}{p_{i',t}^y x_{i'j,t}/p_{j,t}^x x_{j,t}} = \epsilon_j^x \log \frac{z_{ij,t}^x}{z_{i'j,t}^x} + (1-\epsilon_j^x) \log \frac{p_{i,t}^y}{p_{i',t}^y}. \tag{20}$$

so that, for all sectors j, the shares of investment spent on Durable Goods relative to IPP will vary with changes in both their relative prices and relative productivity as inputs, $\frac{z_{ij,t}^x}{z_{i'j,t}^x}$, IBTC in sectoral investment. As with materials, the investment spending shares for the general bundles defined by $\bar{\mathcal{N}}_j^x$ and $\widetilde{\mathcal{N}}_j^x$, $\sum_{i\in\bar{\mathcal{N}}_j^x} \frac{p_{i,t}^y x_{ij,t}}{p_{j,t}^x x_{j,t}}$ and $\sum_{i\in\widetilde{\mathcal{N}}_j^x} \frac{p_{i,t}^y x_{ij,t}}{p_{j,t}^x x_{j,t}}$ respectively, are each time invariant (though not necessarily the same in each sector j) and sum to 1.

Turning our attention to consumption expenditure shares, recall that total consumption expenditures are given by $e_t = \bar{e}_t + \widetilde{e}_t$, where $\bar{e}_t$ and $\widetilde{e}_t$ are expenditures on goods bundled using homothetic and non-homothetic aggregators respectively. Furthermore, motivated by Figure 2, let $s^c = \bar{e}_t/e_t$ represent the constant consumption expenditure share associated with the homothetic bundle combining Construction, Durable Goods, and IPP, and hence $1 - s^c = \widetilde{e}_t/e_t$ is the consumption share associated with the bundle consisting of Nondurable Goods and Services.[12] Then, for consumption goods included in the homothetic bundle, $j \in \bar{\mathcal{N}}^c$, namely Construction, Durable Goods, and IPP, equilibrium consumption shares, $\frac{C_{j,t}}{C_t} = \frac{p_{j,t}^y c_{j,t}}{e_t}$, are given by,

$$\frac{p_{j,t}^y c_{j,t}}{e_t} = \zeta_j s^c,\ j \in \bar{\mathcal{N}}^c, \tag{21}$$

and are, therefore, independent of trend variations in relative prices or income. That is, consumption expenditure shares associated with sectors in $j \in \bar{\mathcal{N}}^c$ will be constant over time. They correspond to the time-invariant rows, θ_j, of the vector $\boldsymbol{\theta}_t$.

For consumption goods included in the non-homothetic bundle, $j \in \widetilde{\mathcal{N}}^c$ (i.e., Nondurable Goods and Services), consumption expenditure shares satisfy

$$\frac{p_{j,t}^y c_{j,t}}{\widetilde{e}_t} = \Theta_{j,t} \left(\frac{p_{j,t}^y}{p_{b,t}^y}\right)^{1-\sigma} \left(\frac{\widetilde{e}_t}{p_{b,t}^y}\right)^{(1-\sigma)(\epsilon_j^c - 1)} \left(\frac{p_{b,t}^y c_{b,t}}{\widetilde{e}_t}\right)^{\epsilon_j^c},\ j \in \widetilde{\mathcal{N}}^c, \tag{22}$$

and

$$\frac{p_{j,t}^y c_{j,t}}{e_t} = \frac{p_{j,t}^y c_{j,t}}{\widetilde{e}_t}(1 - s^c), \tag{23}$$

where recall that b defines a base commodity used to normalize the non-homothetic component of preferences. The choice of the base commodity is arbitrary and, in this case, we set it to our

[12] The non-homotheticity of preferences implies that s^c is a function of not only ρ_t^c but also of the consumption elasticity of expenditures for the non-homothetic bundle, $\mathcal{E}_t$. Maintaining a constant share, s^c, then requires some time variation in ρ_t^c. In equilibrium, this variation turns out to be small.

numéraire, Nondurable Goods. Since $\widetilde{\mathcal{N}}^c$ includes only two sectors, $\frac{p^y_{b,t}c_{b,t}}{\widetilde{e}_t} = 1 - \frac{p^y_{j,t}c_{j,t}}{\widetilde{e}_t}$. That is, low-frequency changes in the consumption share of Nondurable Goods are offset by opposite changes in the share of Services, as suggested by Figure 2, and the combined trends always add up exactly to $1 - s^c$. In equation (22), these changes in turn depend directly on i) exogenous shifts in preferences towards particular goods, $\Theta_{j,t}$, ii) trend variations in relative prices, $\frac{p^y_{j,t}}{p^y_{b,t}}$, and iii) income effects, $\frac{\widetilde{e}_t}{p^y_{b,t}}$, resulting from more fundamental shifts, e.g., in sectoral TFP. Finally, under the maintained assumptions, consumption expenditure shares associated with the bundles defined by $\bar{\mathcal{N}}^c$ and $\widetilde{\mathcal{N}}^c$ satisfy $\sum_{j\in\bar{\mathcal{N}}^c} \frac{p^y_{j,t}c_{j,t}}{e_t} + \sum_{j\in\widetilde{\mathcal{N}}^c} \frac{p^y_{j,t}c_{j,t}}{e_t} = s^c + (1 - s^c) = 1$.

5 Quantifying the Model

To quantify the model, we need to specify two types of objects: fixed model parameters underlying technology and preferences, and time-varying driving processes. We begin by describing the values that we use for the fixed parameters and then present the implied driving processes in the model.

5.1 Parameter Values

We choose parameter values in three ways.

First, the model implies that several of the parameters correspond to nominal expenditure or income shares, and we set the value of these 'share' parameters equal to the sample average share in our dataset. For example, the model-implied value-added share of gross output for sector j is given by the parameter γ_j in equation (4), and we set $\gamma_j = T^{-1}\sum_{t=1}^{T}(V_{j,t}/Y_{j,t})$. The other share parameters are ρ^m_j and ρ^x_j (in equations 5 and 9), ζ^m_{ij} and ζ^x_{ij} (in (7 and 11), α_j (in 8) and ζ^c_j (in 14).

Second, we use the relationships in (19) and (20) to estimate the CES parameters $(\epsilon^m_j, \epsilon^x_j)$ that appear in equations (6) and (10). In particular we estimate the substitution elasticities for intermediate inputs, ϵ^m_j, by regressing the logarithm of relative inputs, $\ln(m_{IPP,j,t}/m_{DUR,j,t})$ onto relative prices, $\ln(p_{IPP,t}/p_{DUR,t})$, using the relative TFP levels, $\ln(z_{IPP,t}/z_{DUR,t})$, as an instrument. Since relative prices are endogenous, our choice of instrument is motivated by the notion that these relative prices respond to the relative value of TFP, which in turn are uncorrelated with the relative value of input biased technical change, $\ln(z^m_{IPPj,t}/z^m_{DUR,j,t})$. The substitution elasticities for investment inputs, ϵ^x_j, are estimated analogously. As described in the online Technical Appendix, we implement this IV estimator using the low-frequency variation in the data using the framework described in Müller and Watson (2017). Importantly, the low-frequency information in the data is limited – it is summarized by $q = 7$ independent observations – and we estimate the parameters using Bayes methods. Table 4 shows production elasticity estimates for each sector with respect to both intermediate inputs and investment goods. Not surprisingly, given the limited low-frequency sample information, there is considerable uncertainly about the values of ϵ^m_j and ϵ^x_j. We use the posterior mean as point estimates and these suggest that, for the production of investment goods, durable goods and IPP are complements in the IPP sector and substitutes in the other sectors. For material inputs, IPP and Durables are complements in the Services and Nondurables sectors and substitutes in the other

Table 4: Estimates of Sectoral Production Elasticities

Sectors	Durables	IPP	Construction	Nondurables	Services
ϵ^x	1.26	0.69	1.43	1.01	1.08
	[0.46,1.14,2.04]	[0.17,0.52,1.20]	[0.37,1.24,2.47]	[0.26,0.84,1.73]	[0.32,0.93,1.81]
ϵ^m	1.69	1.11	2.03	0.86	0.69
	[0.69,1.59,2.61]	[0.37,1.03,1.79]	[1.08,2.01,2.92]	[0.27,0.76,1.43]	[0.14,0.51,1.21]

Notes: The table shows the posterior mean of the elasticities and, in brackets, the 0.17, 0.50, and 0.83 quantiles of the posterior. See the Online Appendix for estimation details.

sectors.

Previous estimation work generally adopt a broader distinction between goods (which include Durables) and services (which include IPP), for example, García-Santana et al. (2021), Herrendorf et al. (2021), or Gaggl et al. (2023).[13] Gaggl et al. (2023) find that goods and services are substitutes in the investment aggregator, while García-Santana et al. (2021) find them to be complements. Because our classification of Durable Goods and IPP as inputs is distinct from these prior exercises, our findings are not directly comparable but they indicate that Durable Goods and IPP can be either substitutes or complements depending on the sector.

The third method for choosing parameter values is informed by previous studies. For σ and ϵ^c (in equation 13) we use values from Comin et al. (2021): $\sigma = 0.26$, $\epsilon^c_{SRV} = 1.65$, and $\epsilon^c_{ND} = 1$ (a normalization).[14] For the values of γ_ℓ and β in (15), we set $\gamma_\ell = 1.0$ and $\beta = 0.96$.[15] Finally, the depreciation rate in each sector, j, is obtained from the relationship that ties that sector's steady state investment share, $\psi^x_j = p^x_j x_j / p^v_j v_j$, to its capital share, $\alpha_j = u_j k_j / p^v_j v_j$, specifically, $\delta_j = \frac{\psi^x_j}{\alpha_j}\left(\frac{1}{\beta} - 1\right) / \left(1 - \frac{\psi^x_j}{\alpha_j}\right)$.

Some of these parameter values are admittedly associated with sizable uncertainty, in particular the input elasticities of substitution, but we find them to be useful benchmarks. The online Technical Appendix summarizes selected results for alternative parameter values.

5.2 Driving Processes

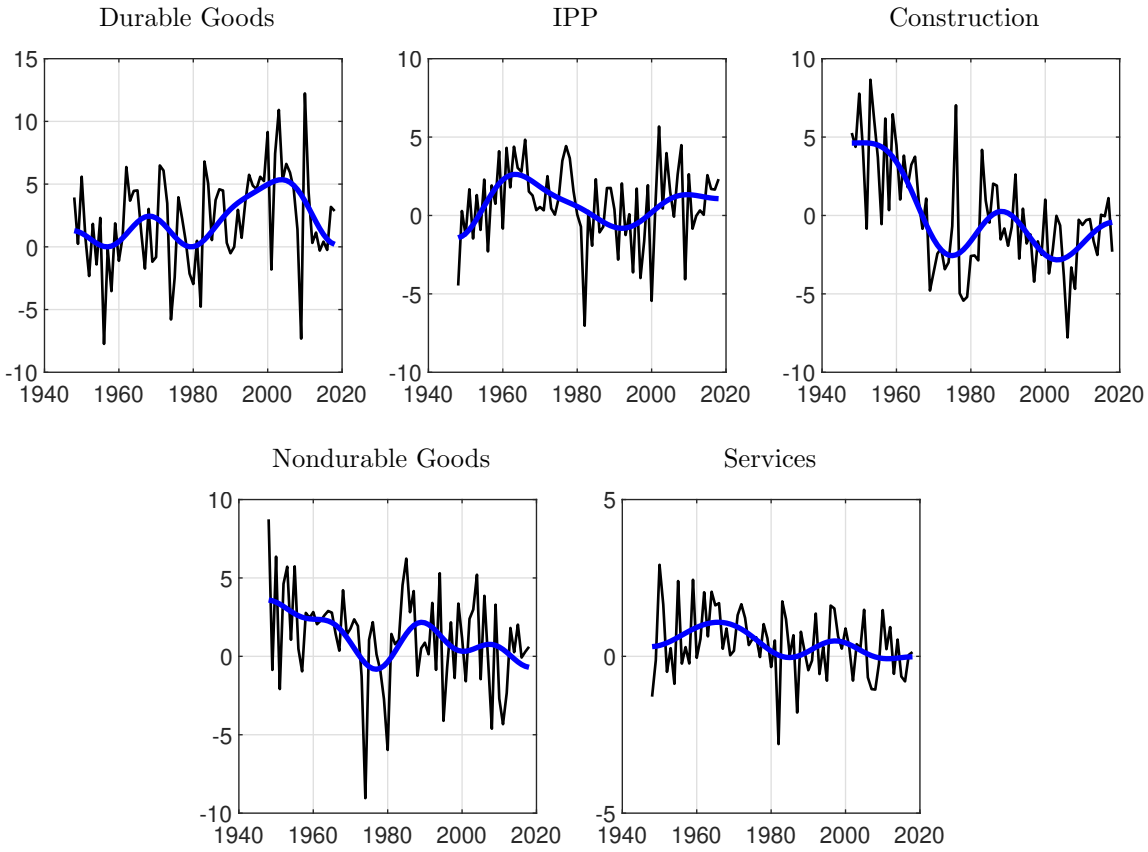
The driving processes include the sector-specific total factor productivity (TFP) processes, $z_{j,t}$, the input-biased technical change processes $(z^m_{ij,t}, z^x_{ij,t})$, the non-homothetic consumption bundle shifters, $\Theta_{j,t}$, the labor supply shifters, $\varphi_{\ell,t}$ and the consumption-share process, ρ^c_t. We discuss these in turn.

[13] In that work, substitution elasticities are estimated using non-linear least squares or GMM, assuming that error terms are uncorrelated with prices.

[14] In principle, $(\sigma, \epsilon^c_{ND})$ can be estimated using low-frequency IV methods analogous to the method used to estimate the CES production parameters $(\epsilon^m_j, \epsilon^x_j)$. However, because of the limited information in the low-frequency time series data, we relied on the Comin et al. (2021) estimates, which are based on micro data.

[15] The Frisch elasticity of labor supply, $1/\gamma_\ell$, varies considerably across studies. Microeconomic estimates tend to be less than 1 while those used to inform aggregate macroeconomic models are typically greater than 1; for surveys see Chetty et al. (2011) and Keane and Rogerson (2012). The value of β is consistent with an interest rate of approximately 4 percent.

Figure 5: Sectoral Total Productivity Growth (Percent)

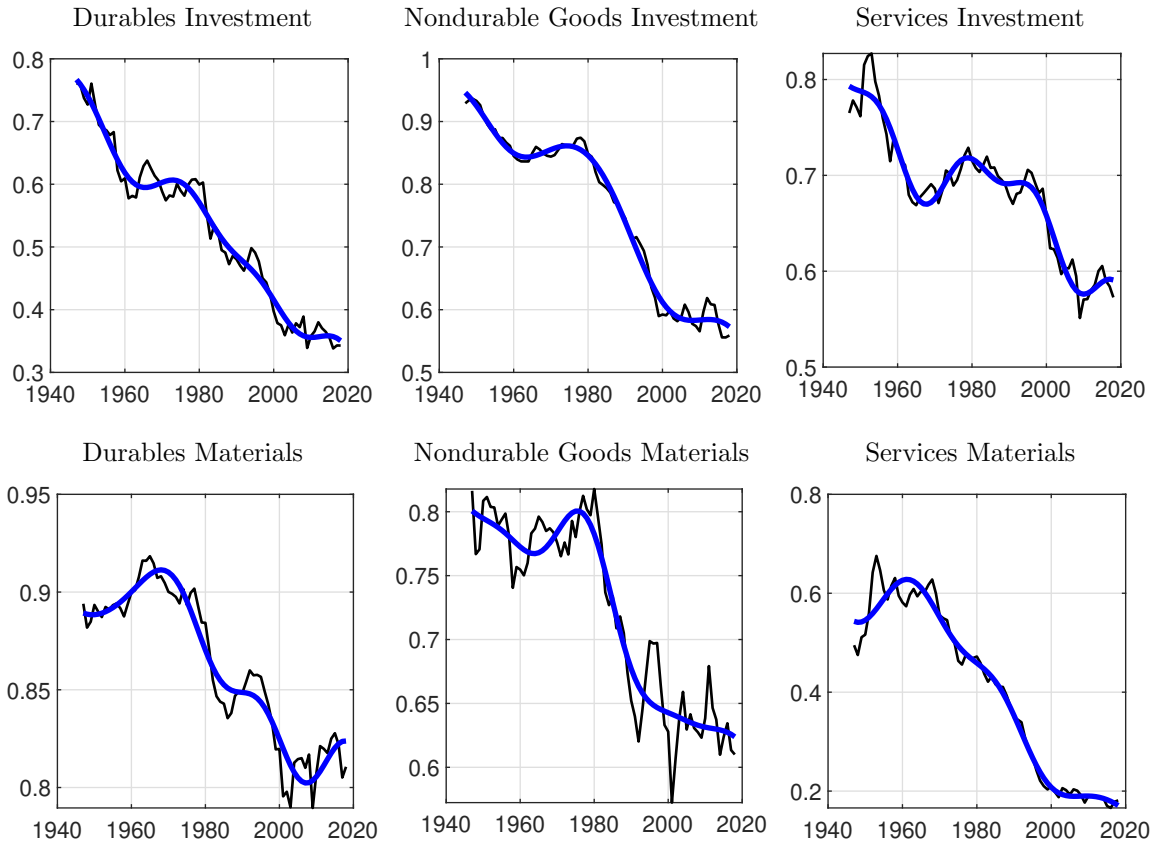


TFP is the Solow residual from equation (8). Figure 5 shows the evolution of TFP growth rates by sector over our sample period, $z_{j,t}$. The panel shows disparate trend variations across sectors. Trend TFP growth in Construction was high in the early part of the sample but then declined and never fully recovered (see, for example, Goolsbee and Syverson (2023), and, for possible explanations, Herkenhoff et al. (2018) or Babalievski et al. (2023)). In contrast, the evolution of trend TFP growth in Durable Goods is virtually the opposite of that in Construction, starting relatively low early in the sample but rising steadily through the end of the 1990s, reflecting in part significant technological improvements in the semiconductor industry during that decade. Trend TFP growth in IPP increases noticeably in the 1950s, and again since the early 1980s, while trend TFP growth in Nondurable Goods starts out relatively robustly early in the sample but gradually falls over time. Finally, often cited as the classic example underlying Baumol's cost disease, trend TFP growth in Services is uniformly low through the entire sample period, with notably smaller trend variations compared to those of other sectors.

The IBTC processes for materials and investment are computed from (19) and (20). Figure 6 illustrates, for the sectors highlighted in Figure 3, the implied evolution of IBTC. Recall that the IBTC processes are normalized so that $z_{DUR,j,t} + z_{IPP,j,t} = 1$, so that the relative productivity of durables is given by $z_{DUR,j,t}/z_{IPP,j,t} = z_{DUR,j,t}/(1-z_{DUR,j,t})$. In all panels, the relative productivity of Durable Goods to IPP inputs falls over time. In the Durable Goods sector for example, new capital

Figure 6: Input-Biased Technical Change

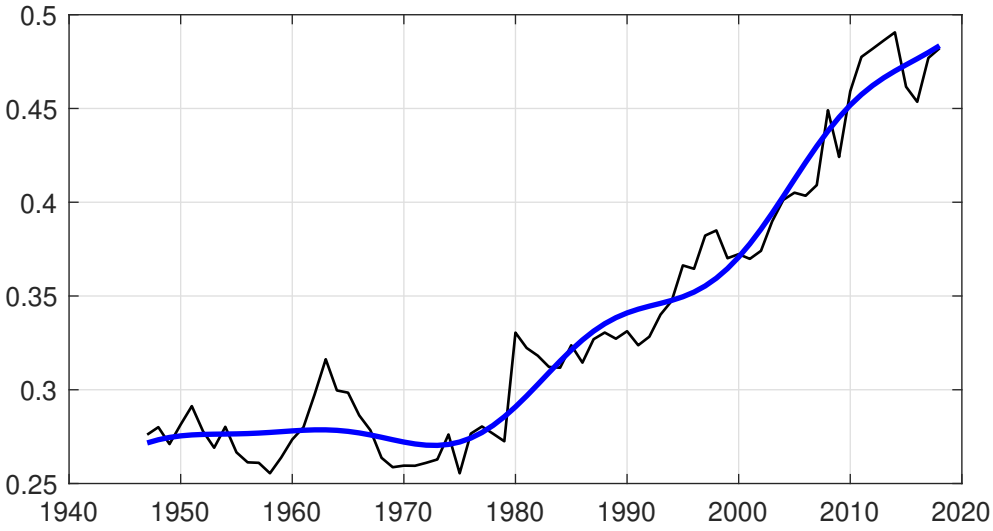
Relative Productivity of Durables Inputs z_{DUR} with $z_{DUR} + z_{IPP} = 1$



made from its own sector is about 4 times as productive as capital produced in the IPP sector at the end of WWII (the top left panel). By 2018, its Durable Goods are only about 40 percent as productive as IPP (0.3/0.7) when employed as new capital in the Durable Goods sector. A similar picture emerges in material input use. In the production of Services, for instance, materials composed of Durable Goods are just about as productive as materials purchased from IPP immediately following WWII (the bottom right panel). By 2018, Durables are only about 1/4 as productive as IPP when used as intermediate inputs in Services (0.2/0.8).

As with IBTC, we can infer exogenous shifts in preferences, $\Theta_{j,t}$, from observations on relative prices, $\frac{p^y_{j,t}}{p^y_{b,t}}$, and consumption expenditures, $\frac{\widetilde{e}_t}{p^y_{b,t}}$, conditional on the preference parameters, σ and ϵ^c_j, equation (22). Since $\Theta_{j,t}$ is only identified up to a base good, b, it represents an exogenous shift in preferences towards goods produced in sector j relative to b. In this case, recall that $\widetilde{\mathcal{N}}^c$ comprises Services and Nondurable Goods, where Nondurables are defined as the base good. Thus, $\Theta_{SVC,t}$ is a measure of exogenous shifts in preferences from Nondurable Goods towards Services. Figure 7 shows this exogenous shift in preferences from Nondurables to Services. From the late 1940s to the mid 1970s, this shift in preferences is relatively flat. It follows that in the context of our model, the decline in the share of Nondurable Goods consumption during that 25-year period, shown in Figure 2, and the corresponding increase in the share of Services consumption over the same period,

Figure 7: Exogenous Shift in Preferences
Services Relative to Nondurable Goods, Θ_{SVC}



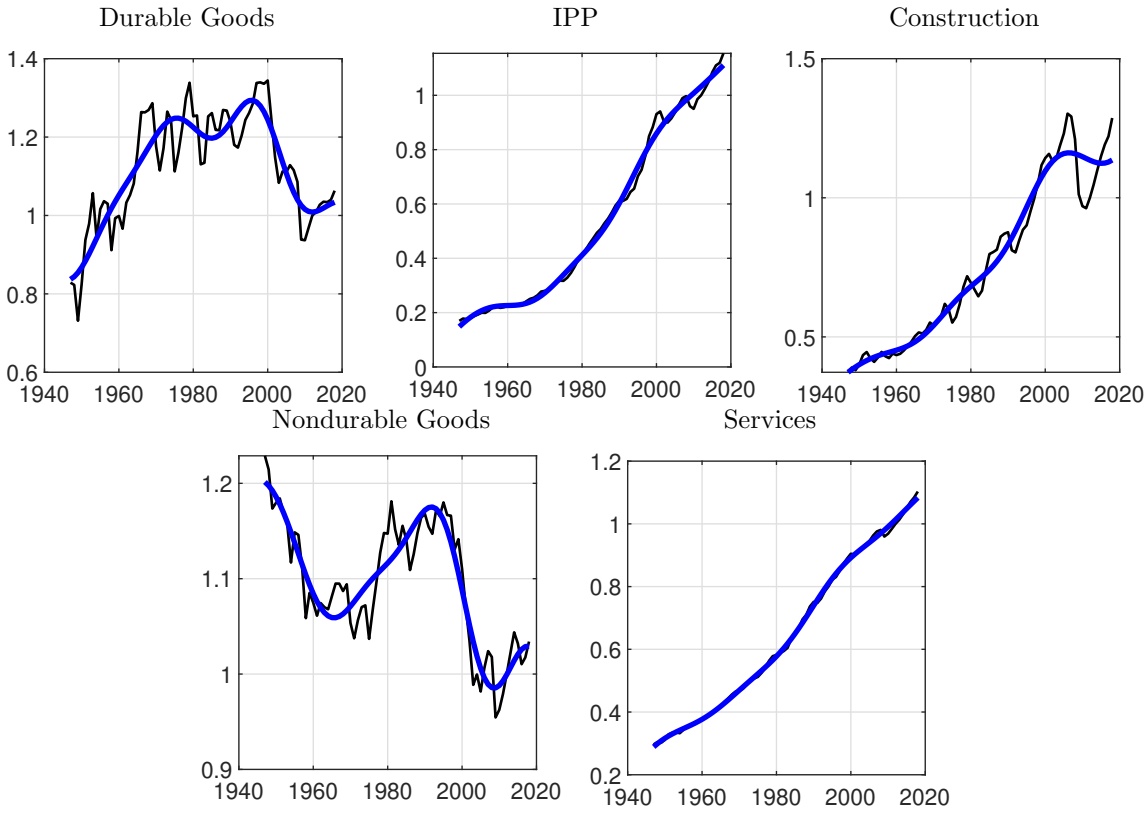
must derive from the general equilibrium effects of other forces on relative prices and consumption expenditures. Beginning in the mid-1970s, preferences steadily shift towards Services and thus begin to directly contribute to their already increasing share and, conversely, the declining consumption share of Nondurable Goods.

The labor supply shifters, $\varphi_{j,t}$, are chosen to match observed sectoral labor input, shown in Figure 8, conditional on changes in labor demand driven by trend variations in sectoral TFP. The measure of labor input produced by KLEMS is quality-adjusted and grows over time in sectors that have gained in importance as a share of GDP, in particular Services and IPP. Labor input is either roughly unchanged or falls over time in sectors that have gradually declined as a share of GDP, namely Durable Goods and Nondurable Goods. Finally, as described above (see footnote 10), the value of ρ_t^c is chosen so that the model-implied combined consumption shares of IPP, Durables, and Consumption are constant and equal to the average share over the sample period. As shown in the online Technical Appendix, variations in ρ_t^c are small.

5.3 Baseline Results

This section presents our baseline results. The key difference with Section 3 is that all sectoral shares in the compositional identity shown in (3), apart from ratios of government spending and net exports to value added, are now endogenous objects. The basic motivation for retaining the assumption of unit-elasticity for various aspects of technology and preferences is twofold. First, it greatly simplifies the analysis and parameterization of the model. Second, the choice of shares we model as constant over time is directly suggested by correspondingly small trend variations in the underlying data. Recall, for example, the consumption shares shown in the lower right-hand panel of Figure 2. In addition, input cost shares in Figure 3 either exhibit little to no trend variation over the sample, for example in the lower left-hand panel, or elsewhere exhibit trend movements in

Figure 8: Labor Input



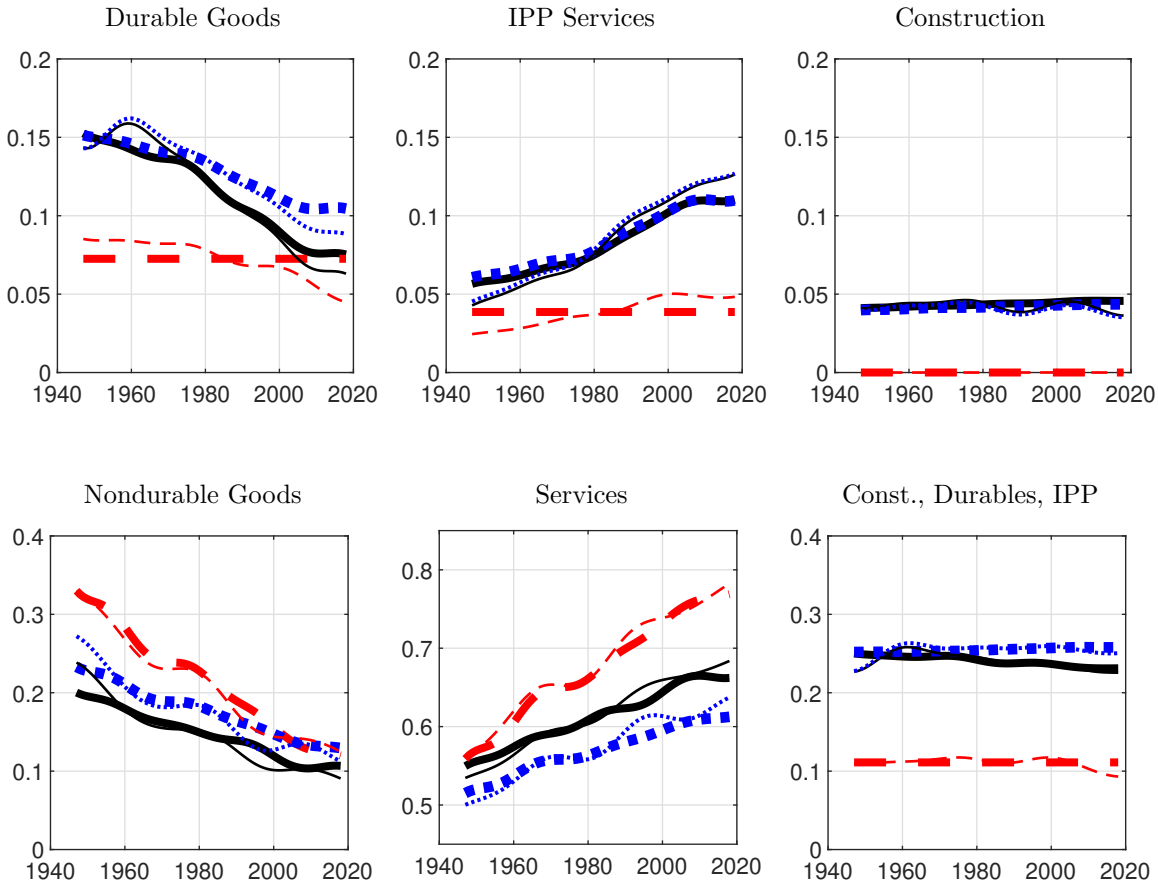
opposite direction with almost exactly offsetting changes.

Figure 9 shows that the baseline structural model, driven by sectoral TFP, IBTC, and preference shifts, is able to capture essentially all of the salient trend variations in sectoral value added shares of GDP, as well as those of the constituent subshares that make up the compositional identity (3). Since the model treats consumption shares in Durable Goods and IPP as exactly constant over time, it misses some of the secular changes in these shares, and thus understates slightly trend changes in these sectors' corresponding value added shares of GDP. However, the structural model essentially captures all of the trend changes (or structural change) in Nondurable Goods and Services. In addition, in the lower right-hand panel of Figure 9, the trend consumption and value added shares combining sectors employed mainly in the production of inputs for the U.S. economy, namely Construction, Durable Goods, and IPP, are largely constant over time.

Exploring further the underlying constituents of structural change, Figure 10 shows the implications of the model for the input cost shares associated with investment and materials composed of Durable Goods and IPP in various sectors. By and large, IPP input shares gradually increase over time while those associated with Durable Goods decline, except for the Durables Goods sector, where intermediate input shares are approximately constant throughout the sample. Furthermore, recall that our assumptions on technology imply that trend variations in input shares of IPP and Durable Goods, whether used as investment or materials, always offset each other and sum to a constant, $\sum_{i\in\widetilde{\mathcal{N}}_j^m} \frac{p_{i,t}^y m_{ij,t}}{p_{j,t}^m m_{j,t}} = \rho_j^m$, $\sum_{i\in\widetilde{\mathcal{N}}_j^x} \frac{p_{i,t}^y x_{ij,t}}{p_{j,t}^x x_{j,t}} = \rho_j^x$, which indeed conforms to the data well.

Figure 9: Structural Change: Model vs. Data

Consumption Only (red dashed), with Production Linkages (blue dots), Value Added Shares (solid black)



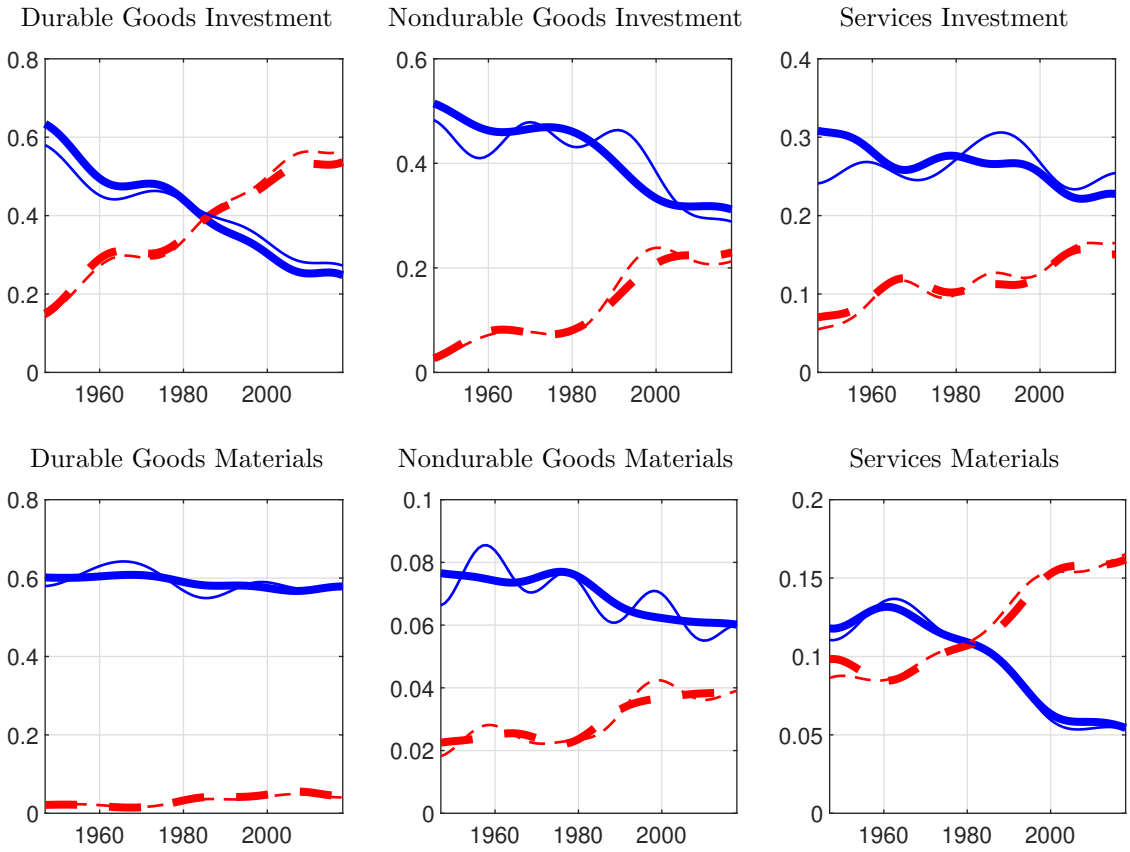
Notes: The red lines describe the low frequency evolution of consumption expenditure shares, $\boldsymbol{\theta}_t$. The blue lines add the contributions from production linkages to value added shares in GDP, $\boldsymbol{\Gamma}_{d,t}$, $\boldsymbol{\Phi}_t$ $\boldsymbol{\Omega}_t$, and $\boldsymbol{\Psi}_t^x$. Sectoral value added shares in GDP are shown in black. Thin lines depict data trends and thick lines depict trends implied by the structural model.

Of course, the structural changes depicted in Figures 9 and 10 reflect, in equilibrium, trend variations in underlying fundamentals such as TFP, IBTC or exogenous shifts in preferences, in a way that is not immediately obvious from these baseline figures. Thus, the next section addresses the structural decomposition of secular changes in the composition of the U.S. economy in terms of these underlying forces.

6 Structural Compositional Accounting and Forecasting

In this section, we revisit the notions of compositional accounting and compositional forecasting introduced in Section 3, but from the standpoint of the fundamental driving forces and mechanisms described in Sections 4 and 5. In particular, we are interested in how low-frequency changes in economic fundamentals, both on the supply side, through productivity, and on the demand side, through

Figure 10: Trend Input Cost Shares: Model vs. Data
Durable Goods (blue) vs. IPP (red), Selected Sectors



Notes: In the top row, the blue and red lines describe the low frequency evolution of investment spending on Durable Goods and IPP, respectively, as a share of total investment spending for the different sectors shown. The bottom row depicts the analogous shares for spending on materials.. Thin lines depict data trends and thick lines depict trends implied by the structural model.

exogenous shifts in preferences, including labor supply, have contributed to structural change. We further explore what these low-frequency trends imply for trend forecasts of the composition of the U.S. economy.

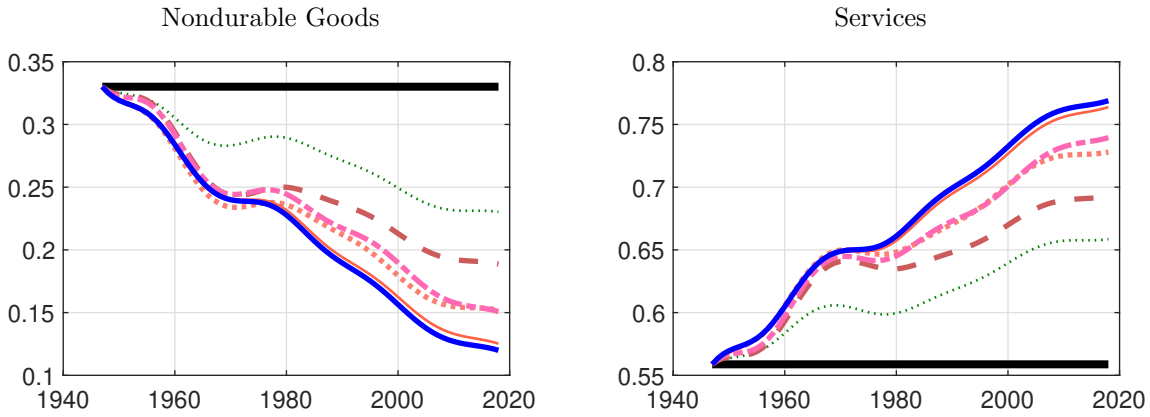
6.1 Historical Structural Accounting

Figure 11 illustrates the cumulative contributions from trend variations in economic fundamentals (i.e., TFP, IBTC, etc.) to the gradual fall and rise, respectively, in the consumption expenditure shares of Nondurable Goods and Services (recall that consumption shares in the other sectors are roughly constant).

In Figure 11, the flat black lines depict the shares of Nondurables and Services in total consumption expenditures at the start of the sample, and thus do not change in the absence of any shocks. Interestingly, in equilibrium, far and away the most important fundamentals accounting for the secular evolution of these expenditure shares arise from the supply side through trend variations

Figure 11: Cumulative Decomposition of Consumption Shares

TFP with Homothetic Preferences (thin dots), TFP with Non-Homothetic Preferences (thick dots), with IBTC (dash), with Preference Shifts (dash-dot), with Labor Supply Shifts (thin), with Government and Net Exports (thick)



in TFP, $z_{j,t}$, accompanied by positive income effects for Services.

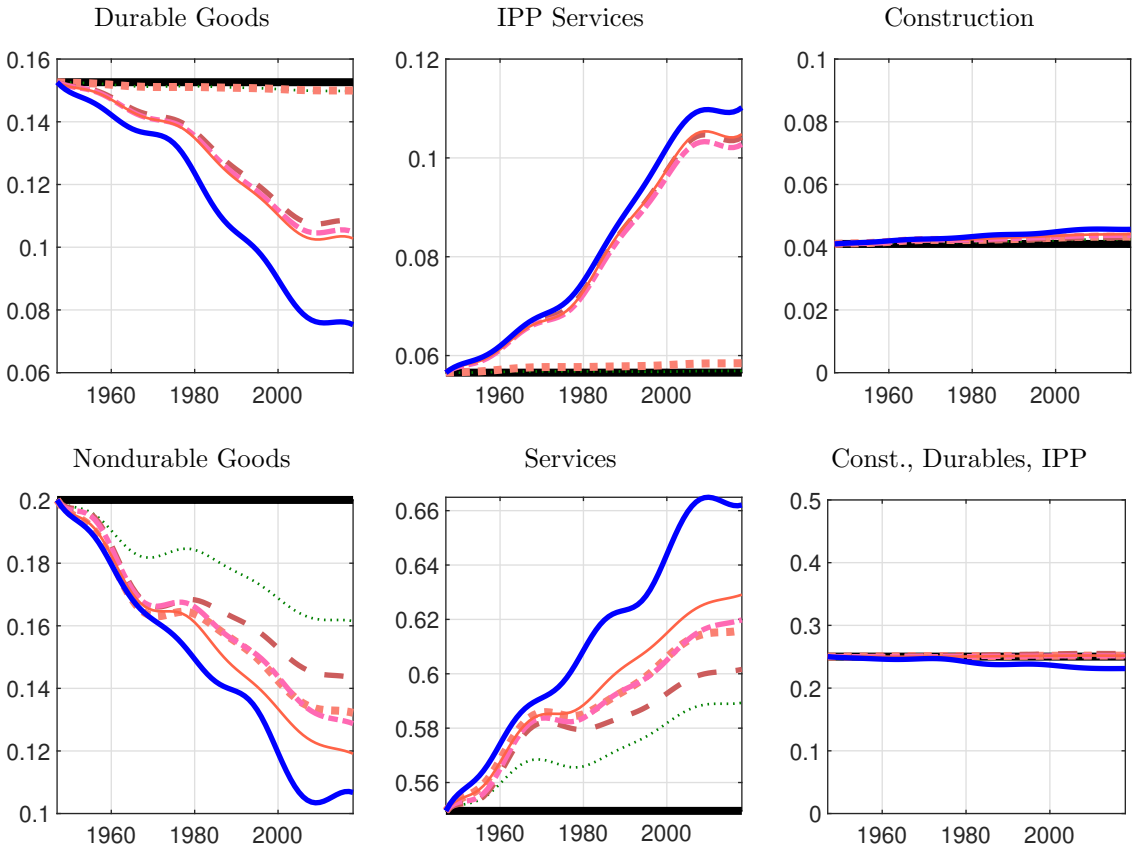
To gain insight into the role of sectoral TFP for the changing composition of consumption expenditures, observe in Figure 5 that TFP growth in Services is almost uniformly lower than in any other sector throughout the post-war period, including Nondurable Goods. This feature of sectoral TFP growth is the reason why Services are often used to illustrate Baumol's cost disease, and implies a rising price of Services relative to Nondurables over time. Suppose we set the expenditure elasticity for Services, ϵ^c_{SVC}, to 1 so that the aggregator function used to bundle the consumption of Nondurables and Services, equation (12), becomes standard CES. Then, since Services and Nondurable Goods are complements in consumption, $\sigma < 1$, an increasing price of Services relative to Nondurables (driven by low productivity growth) implies an increasing relative expenditure share of Services. The line labeled 'TFP with Homothetic Preferences' in Figure 11 plots the relative price effect on consumption shares arising from this counterfactual. However, as TFP grows over time across all sectors, households become richer and total consumption expenditures increase. In our baseline with non-homothetic preferences, $\epsilon^c_{SVC} > 1$, the income effect associated with rising expenditures, $\widetilde{e}_t/p_{b,t}$, in equation (22), adds additional impetus to the consumption of Services, such that the share of Services increases by twice as much, the lines labeled 'TFP with Non-Homothetic Preferences' in Figure 11.

Trend changes in sectoral TFP, and their attendant effects on relative prices and income, explain essentially all of the secular shifts in the consumption shares of Nondurables and Services until about 1980 (the dotted lines). This result is consistent with the relatively flat trend in preferences for Services relative to Nondurables seen in Figure 7. After 1980, exogenous shifts in preferences favoring Services, $\Theta_{SVC,t}$, play somewhat of a role in explaining the compositional shift in consumption from Nondurables to Services, but at the same time are counteracted by the effects of IBTC.

Beyond consumption expenditure shares, Figure 12 illustrates the structural accounting of historical trends in sectoral shares of GDP, i.e., the structural accounting of structural change. The central lesson from the figure is that no one underlying driving process dominates structural change,

Figure 12: Cumulative Decomposition of Value Added Shares

TFP with Homothetic Preferences (thin dots), TFP with Non-homothetic Preferences (thick dots), with IBTC (dash), with Preference Shifts (dash-dot), with Labor Supply Shifts (thin), with Government and Net Exports (thick)



but rather that structural change derives from different forces and mechanisms in different sectors. In particular, consumption income effects barely spill over to the investment sectors.

Consistent with the data, the value added share of Construction never changes very much. Its share of total consumption expenditures is treated as constant, and so are the input shares corresponding to the purchase of Construction materials or investment goods in every sector. In producing structures, Construction does employ Durable Goods and IPP with time-varying input shares, but this attribute of production remains relatively immaterial.

Figure 12 shows that relatively little of the trend decline in the share of Durable Goods of GDP stems from TFP. In contrast, with the efficiency of this sector's goods gradually falling relative to IPP when used as inputs in different sectors (Figure 6), IBTC accounts for approximately 1/2 of the decline in its importance to the U.S. economy throughout the sample. Other forces, such as exogenous shifts in preferences or labor supply, do little to affect structural change in Durable Goods. Instead, the remaining 1/2 of the trend decline in that sector's share of GDP is accounted for by net exports, or the fact that the U.S. has gradually shifted from being an exporter to being an importer of durable goods over time.

In the IPP sector, IBTC essentially accounts for all of the trend increase in value added relative to

GDP. In particular, as with Durable Goods, sectoral TFP plays a relatively minor though not entirely negligible role. Because IPP acts primarily as a provider of intermediate inputs and investment to other sectors, demand forces in the form of preference shifts, either in consumption or labor supply, play essentially no part in explaining its increasing share of GDP over the post-WWII period. Furthermore, income effects arising from non-homothetic preferences do not play much of a role for IPP, Durable Goods, or Construction.

Finally, in Figure 12, sectoral TFP accounts for about 1/2 the trend evolution in the shares of Nondurable Goods and Services of GDP at the end of the sample, and nearly all of this evolution until around 1980, with about equal contributions from relative price and income effects. These findings essentially mimic those related to their consumption expenditure shares shown in Figure 11. Recall that in the compositional identity (3), the trend behavior of the value added shares of Nondurables and Services arose largely from the trend behavior of their consumption expenditure shares. It is not surprising, therefore, that the driving processes that explain the secular compositional shift in consumption from Nondurables to Services would also explain their compositional shift in value added. That said, one difference between Figure 11 and Figure 12 is that shifts in labor supply, $\varphi_{j,t}$, explain secular changes in the GDP shares of Nondurables and Services to a greater degree (in Figure 12,) than their corresponding consumption shares (in Figure 11). Shifts in labor supply directly affect sectoral labor input and, therefore, production and value added while only indirectly acting on consumption through their influence on relative prices.

6.2 Structural Compositional Forecasting

At the end of Section 3, we applied reduced-form low frequency prediction methods to forecast 20-year ahead trend value-added shares. In this section we use the structural model to construct analogous forecasts. The structural forecasts differ from the reduced form forecasts in two ways. First, they are based on forecasts for each of the exogenous driving processes underlying the structural model, TFP, IBTC, exogenous preference shifts, etc. These driving processes are assumed to be independent of each other so, for example, TFP is independent of preference shifts. Second, while reduced-form forecasts used a logit transformation to enforce the constraint that the value-added shares sum to unity, the structural model endogenously enforces this constraint for the structural forecasts.

To compute the structural forecasts, we first use the methods described in Section 3.2.1 to construct predictive distributions for the trend values of the exogenous variables in the structural model. We do this one driving process at a time, first TFP, then IBTC in materials, then IBTC in investment, etc. Although the driving prcesses are assumed to be independent of each other, we take into account the co-movement across sectors within a driving process, so TFP in Durable Goods potentially comoves with TFP in IPP and so on. Then we compute the implied predictive distribution for equilibrium value-added shares in the model. The details underlying theses calculation are provided in the online Technical Appendix. We note that the forecasts ignore uncertainty in the value of the model parameters described in Section 5.1. This approximation, while not ideal, simplifies the analysis considerably albeit at the cost of underestimating the uncertainty in the forecasts. Table 5 summarizes the resulting forecasts.

Table 5: Forecasted Changes in Trend Sectoral Shares of GDP
2018 to 2038 (Percentage Points)

Sectors	Durables	IPP	Construction	Nondurables	Services
Historical 20 Year Avg	-2.2	2.3	-0.1	-4.1	4.1
Reduced Form	-1.5	3.5	-0.3	-2.6	0.9
Forecast 2018-2038	[-2.4,-0.4]	[2.4,4.5]	[-1.1,0.8]	[-3.6,-1.4]	[-1.2,2.5]
Structural Model	-1.4	1.5	0.0	-1.3	1.2
Forecast 2018-2038	[-2.1,-0.8]	[0.6,2.4]	[0.0,0.1]	[-2.0,-0.6]	[0.1,2.2]
By Process:					
TFP ($z_{j,t}$)	0.0	0.0	0.0	-0.7	0.7
IBTC ($z^m_{ij,t}, z^x_{ij,t}$)	-0.9	1.4	0.0	-0.2	-0.2
Preferences ($\Theta_{j,t}$)	-0.1	0.0	0.0	-0.3	0.4
Labor Supply ($\varphi_{j,t}$)	-0.1	0.1	0.0	-0.3	0.3
G and NX ($\Psi_{j,t}$)	-0.5	0.2	0.0	-0.3	0.6

The first 3 rows of Table 5 compare structural forecasts of 20-year changes in trend sectoral shares with those from the reduced-form forecasts in Section 3 as well as historical 20-year averages. Similar to reduced-form forecasts, the structural forecasts indicate a slowdown in structural change. For example, the trend share of Nondurables fell on average by 4.1 percentage points historically every 20 year period, but is forecasted to decline by 2.6 percent over the next 20 years in Section 3, and by only 1.3 percent in the structural model.

While both reduced-form and structural forecasts indicate a slowdown in the pace of structural change, a key difference between the two sets of forecasts lies in their relative magnitudes. For example, expected changes in the relative size of Services and Nondurables in the reduced-form forecasts do not offset each other as they have historically, and neither do expected changes in the relative size of IPP and Durables. In contrast, structural forecasts indicate that changes in sectoral value added shares are expected to abate in a way that largely respect historical patterns. Thus, the IPP share of GDP is expected to rise by 1.5 percentage points over the next 20 years, largely offset by a decline in the relative size of Durables by 1.4 percentage points. A similar pattern emerges between Nondurables and Services.

For additional insight into the differences between reduced-form and structural forecasts, Table 5 also decomposes the structural forecasts of changes in trend sectoral shares by driving process. That is, we construct forecasts of trend sectoral shares of GDP conditional on trend forecasts of each of the driving processes one at a time.[16] Consistent with our historical findings, forecasts of trend

[16]Because the model is nonlinear, the by-process point forecasts do not sum to the structural model forecast.

changes in the shares of Nondurables and Services are dominated by TFP, but consistent with the historical experience shown Figure 12, the structural forecasts associated with TFP are such that expected changes in the consumption shares of Nondurables and Services will largely offset each other. In Durable Goods and IPP, forecasts of IBTC mostly govern the forecasts of changes in trend value added shares, subject to their corresponding input shares largely offsetting each other as in Figure 10. Construction's share of GDP, consistent with the general history of that sector, is not forecasted to change materially.

7 Concluding Remarks

We explore historical shifts in sectoral shares of GDP both from a compositional accounting standpoint, where the evolution of these shares reflect underlying changes in the observed composition of consumption and production, and within an economic model, where structural change emerges endogenously in general equilibrium and is driven by economic fundamentals. In a compositional accounting sense, the steady shift from goods to services that characterizes the post WWII U.S. economy was primarily driven by two key changes. First, Intellectual Property Products gradually became the dominant source of inputs in U.S. production at the expense of Durable Goods. Second, the consumption share of Services steadily increased in a way that was almost exactly offset by a decline in the consumption share of Nondurable Goods. Our economic model then indicates that the steady increase in IPP's share of GDP stemmed mostly from increases in the relative efficiency of IPP inputs in production, reflecting input-biased technical change. At the same time, trends in sectoral TFP growth, and their attendant effects of relative prices and income, were the main driver underlying compositional shifts in consumption.

Beyond historical shifts in the composition of U.S. GDP, we further tackle the question of how to forecast compositional trends going forward. First, we construct reduced-form compositional forecasts. These require an ad-hoc transformation of the data to ensure that forecasts of trend sectoral shares satisfy their adding-up constraint. Second, we present trend compositional forecasts constructed from forecasts of economic fundamentals within a model where the enforcement of adding-up constraints on sectoral shares is endogenous. Both reduced-form and structural forecasts project historical trends to continue over the next two decades, but at reduced rates, indicating a slowdown in the pace of structural change.

References

Aitchison, J. (1986). *The Statistical Analysis of Compositional Data*. Chapman and Hall: New York.

Babalievski, F., K. Herkenhoff, L. Ohanian, and E. Prescott (2023). The Impact of Commercial Real Estate Regulations on U.S. Output. NBER Working Paper 31895.

Boppart, T., P. Kiernan, P. Krusell, and H. Malmberg (2023). The macroeconomics of intensive agriculture. NBER Working Papers 31101, National Bureau of Economic Research, Inc.

Casal, L. and J. Caunedo (2024). On the investment network and development. mimeo.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011, May). Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review 101*(3), 471–75.

Choi, J. and A. T. Foerster (2017). The Changing Input-Output Network Structure of the U.S. Economy. *Economic Review 102*(2), 23–49.

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

Foerster, A., A. Hornstein, P.-D. Sarte, and M. Watson (2022). Aggregate implications of changing sectoral trends. *Journal of Political Economy 130*, 3286–3333.

Gaggl, P., A. Gorry, and C. Vom Lehn (2023). Structural change in production networks and economic growth. CESifo Working Paper 10460.

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

Goolsbee, A. and C. Syverson (2023). The Strange and Awful Path of Productivity in the U.S. Production Sector. NBER Working Paper 30845.

Greenwood, J., Z. Hercowitz, and P. Krusell (1997). Long-Run Implications of Investment-Specific Technological Change. *American Economic Review 87*(3), 342–362.

Herkenhoff, K., L. Ohanian, and E. Prescott (2018). Tarnishing the Golden and Empire States: Land-Use Regulations and the U.S. Economic Slowdown. *Journal of Monetary Economics 93*, 89 – 109.

Herrendorf, B., R. Rogerson, and A. Valentinyi (2013). Two perspectives on preferences and structural transformation. *American Economic Review 103*(7), 2752–89.

Herrendorf, B., R. Rogerson, and A. Valentinyi (2014). Growth and structural transformation. *Handbook of economic growth 2*, 855–941.

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

Keane, M. and R. Rogerson (2012, June). Micro and macro labor supply elasticities: A reassessment of conventional wisdom. *Journal of Economic Literature 50*(2), 464–76.

Müller, U. K. and M. W. Watson (2008). Testing Models of Low-Frequency Variability. *Econometrica 76*, 979–1016.

Müller, U. K. and M. W. Watson (2016). Measuring Uncertainty about Long-Run Predictions. *Review of Economic Studies 83*, 1711–1740.

Müller, U. K. and M. W. Watson (2017). Low-frequency econometrics. In B. Honoré and L. Samuelson (Eds.), *Advances in Economics: Eleventh World Congress of the Econometric Society*, Volume II, pp. 63–94. Cambridge University Press.

Müller, U. K. and M. W. Watson (2018). Long-Run Covariability. *Econometrica 86*(3), 775–804.

Müller, U. K. and M. W. Watson (2020). Low-Frequency Analysis of Economic Time Series, *in preparation*. In *Handbook of Econometrics*. Elsevier.

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

Restuccia, D. and R. Rogerson (2017). The causes and costs of misallocation. *Journal of Economic Perspectives 31*(3), 151–74.

vom Lehn, C. and T. Winberry (2021). The Investment Network, Sectoral Comovement, and the Changing U.S. Business Cycle. *The Quarterly Journal of Economics 137*(1), 387–433.

Appendix A Defining Trends

This appendix expands the discussion of Section 2.3 about computation of the low-frequency trends. Using the notation introduced there, let x_t denote a time series that is available from $t = 1$ through $t = T$. The low-frequency trend, denoted $\hat{x}_t$, is the fitted value from the regression of x_t onto a constant, possibly a time trend, and a set of low-frequency periodic functions. The details are provided here for the regression that includes a constant but excludes the time trend. The online Technical Appendix discusses the trend case.

To begin, consider the cosine function $\Psi_j(s) = \sqrt{2}\cos(sj\pi)$ on $s \in [0,1]$ with period $2/j$, and let $\mathbf{\Psi}(s) = [\Psi_1(s), \Psi_2(s), ..., \Psi_q(s)]'$ denote a vector of these functions with periods 2 through $2/q$. The regression uses the matrix of periodic regressors given by $\mathbf{\Psi}_T$, a $T \times q$ matrix with t^{th} row $\mathbf{\Psi}\left((t-1/2)/T\right)'$. Note that these regressors capture periodicities longer than $2T/q$. Including a $T \times 1$ vector of 1s, $\mathbf{1}_T$, for the constant term, the matrix of regressors is then given by $\mathbf{\Psi}_T^0 = [\mathbf{1}_T, \mathbf{\Psi}_T]$.

Let $x_{1:T} = (x_1, x_2, ..., x_T)'$ denote the $T \times 1$ vector of observations. The fitted values from the regression of $x_{1:T}$ onto $\mathbf{\Psi}_T^0$ is $\hat{x}_{1:T} = \mathbf{\Psi}_T^0 \mathbf{X}_T^0$, where $\mathbf{X}_T^0 = (\mathbf{\Psi}_T^{0\prime}\mathbf{\Psi}_T^0)^{-1}\mathbf{\Psi}_T^{0\prime}x_{1:T}$ are the OLS regression coefficients. It turns out that this expression can be simplified because the specific form used for the cosine weights implies that the columns of $\mathbf{\Psi}_T^0$ are orthogonal with $T^{-1}(\mathbf{\Psi}_T^{0\prime}\mathbf{\Psi}_T^0) = \mathbf{I}_{q+1}$; this yields the simplified formula $\mathbf{X}_T^0 = T^{-1}\mathbf{\Psi}_T^{0\prime}x_{1:T}$.

Notice that the first element of $\mathbf{X}_T^0$ is the sample mean, $\overline{x}$, and the other elements are given by $\mathbf{X}_{j+1,T}^0 = T^{-1}\sum_{t=1}^T \Psi_j((t-0.5)/T)x_t$, which are low-frequency weighted averages of x_t (these weighted averages are sometimes referred to as the 'cosine transforms' of $x_{1:T}$). Because the elements of $\mathbf{X}_T^0$ represent averages of the T observations, they are (approximately) normally distributed and readily analyzed using standard statistical methods (see Müller and Watson (2020) and the online Technical Appendix).

In our application $T = 72$ years and we set $q = 7$. Thus, the fitted values $\hat{x}_{1:T}$ capture the low-frequency variability in $x_{1:T}$ corresponding to periods longer than $2T/q \approx 20$ years.

Appendix B Trend Compositional Forecasting

This appendix expands the discussion of Section 3.2.1 about forecasting long-run trends. The appendix considers a scalar random variable; the extension for vector valued random variables is straightforward and is included in the online Technical Appendix.

Using the notation in Section 3.2.1, let $t = 1, ..., T_{IS}$ denote the 'in-sample period' and $t = 1, ..., T_{FS}$ denote the 'full-sample period', were $T_{FS} >> T_{IS}$. Interest focuses on the value of the low-frequency trend constructed from the full-sample. Using the notation from Appendix A, these values are $\hat{x}_{1:T_{FS}} = \mathbf{\Psi}_{T_{FS}}^0 \mathbf{X}_{T_{FS}}^0$. Because $\mathbf{\Psi}_{T_{FS}}^0$ is deterministic, forecasts of $\hat{x}_{1:T_{FS}}$ only require forecasts of the OLS coefficients $\mathbf{X}_{T_{FS}}^0$.

We construct forecasts of $\mathbf{X}_{T_{FS}}^0$ using the low-frequency information from the in-sample observations that are summarized in $\mathbf{X}_{T_{IS}}^0$. Consistent with the intuition provided in Appendix A, results in Müller and Watson (2016) and Müller and Watson (2020) imply that under a fairly general as-

sumptions about the stochastic process describing x_t, the in-sample and full-sample values of $\mathbf{X}^0_T$ are (approximately) jointly normally distributed with

$$\begin{bmatrix} \mathbf{X}^0_{T_{IS}} \\ \mathbf{X}^0_{T_{FS}} \end{bmatrix} \overset{a}{\sim} \mathcal{N}\left(0, \sigma^2_{LR} \begin{bmatrix} \Omega_{11} & \Omega_{12} \\ \Omega_{21} & \Omega_{22} \end{bmatrix}\right) \tag{B.1}$$

where σ^2_{LR} is the long-run variance for the stochastic process describing x_t, and the matrix Ω depends on the form of this stochastic process, the cosine weights, as well as T_{IS} and T_{FS} (see the online Technical Appendix for the details of these expressions). From (B.1), the usual conditional normal formula yields

$$\mathbf{X}^0_{T_{FS}} | \mathbf{X}^0_{T_{IS}} \overset{a}{\sim} \mathcal{N}\left(\Omega_{21}\Omega_{11}^{-1}\mathbf{X}^0_{T_{IS}}, \sigma^2_{LR}(\Omega_{22} - \Omega_{21}\Omega_{11}^{-1}\Omega_{12})\right). \tag{B.2}$$

In the forecasting applications in this paper, the matrix Ω is computed under the assumption that x_t follows an $I(1)$ process with a diffuse Gaussian prior for its initial condition (and drift, if present) while the long-run variance, σ^2_{LR}, is estimated using a Newey-West estimator. Forecasts of $\mathbf{X}^0_{T_{FS}}$, and their prediction intervals, can then be readily obtained using draws from the conditional distribution (B.2). These in turn yield forecasts and prediction intervals for the long-run trends $\hat{x}_{1:T_{FS}} = \mathbf{\Psi}^0_{T_{FS}}\mathbf{X}^0_{T_{FS}}$.