



HCEO WORKING PAPER SERIES

Working Paper



HUMAN CAPITAL AND
ECONOMIC OPPORTUNITY
GLOBAL WORKING GROUP

The University of Chicago
1126 E. 59th Street Box 107
Chicago IL 60637

www.hceconomics.org

INCOME-DRIVEN LABOR-MARKET POLARIZATION*

Diego Comin Ana Danieli Martí Mestieri
Dartmouth College Northwestern FRB Chicago & Northwestern

June 27, 2020

Abstract

We propose a mechanism for labor-market polarization based on the nonhomotheticity of demand that we call the income-driven channel. Our mechanism builds on a novel empirical fact: expenditure elasticities and production intensities in low- and high-skill occupations are positively correlated across sectors. Thus, as income grows, demand shifts towards expenditure-elastic sectors, and the relative demand for low- and high-skill occupations increases, causing labor-market polarization. A calibrated general-equilibrium model suggests this mechanism accounts for 90% and 35% of the increase in the wage-bill share of low- and high-skill occupations observed in the US during 1980-2016, and for 64% and 28% of the rise in the employment shares of low- and high-skill occupations. This mechanism is similarly important for the polarization of labor markets in Western Europe during 1980-2016, as well as in the US during earlier decades and, possibly, the near future.

Keywords: Labor-market polarization, Nonhomothetic Demand.

JEL Classification: E21, E23, J23, J31.

*We thank Richard Rogerson and Christian Siegel for insightful discussions of the paper. We also thank Gadi Barlevy, David Dorn, Matthias Doepke, Thibault Fally, Koichiro Ito, Ryan Kellogg, Kiminori Matsuyama, Rachel Ngai, Yongs Shin, Michael Sposi, Chris Tonetti, Joseph Zeira, and attendees at various seminars and conferences for helpful comments and discussions. Comin expresses his gratitude for the financial support from the European Commission under the FRAME project and from the NSF. Mestieri is grateful to CREI and the Chicago Fed for their hospitality as he conducted part of this research at their locations. All errors are our own. Contact: marti.mestieri@northwestern.edu

1 Introduction

There has been a polarization of labor market outcomes in the US and Western Europe since the 1980s. For example, the US wage bill of high- and low-skill workers relative to that of middle-skill workers has more than doubled since 1980.¹ This polarization of wage bills reflects the polarization of both hours worked and wages across occupations. The social significance of these trends has given rise to a vast literature on the drivers of polarization that has focused on technical change, automation, trade and offshoring, and de-unionization, among others.

In this paper, we propose and quantify a new channel to explain labor-market polarization. Through our mechanism, labor market outcomes become increasingly polarized as aggregate household income rises. Consequently, we name this mechanism the income-driven channel for polarization. The income-driven channel builds on a novel empirical fact: expenditure-elastic sectors intensively use high- and low-skill occupations in production, whereas there is no such correlation for middle-skill occupations. As income grows, demand shifts toward sectors with higher expenditure elasticity. Since these sectors are intensive in high- and low-skill occupations, the reallocation of sectoral demand causes an increase in the relative demand for high- and low-skill occupations. This leads to a hollowing out of middle-skill occupations and the polarization of workers' earnings.²

Using household-level data, we document significant variation in expenditure elasticities across sectors.³ Figure 1 shows that sectors with high-expenditure elasticities grew at a higher rate than the US average between 1980 and 2016. The correlation between sectoral expenditure elasticity and sectoral value-added growth is over 0.8. Central to our theory, Figures 2a and 2b show that there is a strong positive correlation between the expenditure elasticity of a sector and its intensity in low- and high-skill occupations, measured as total payments to an occupation over sectoral value added.⁴ In contrast, there is a mild negative relationship with the intensity in middle-skill occupations (Figure 2c). The corresponding correlations are 0.93 and 0.82 for high- and low-skill occupations and -0.04 for middle-skill occupations. Moreover, Figure 2d shows that the distribution of factor intensities across sectors is stable

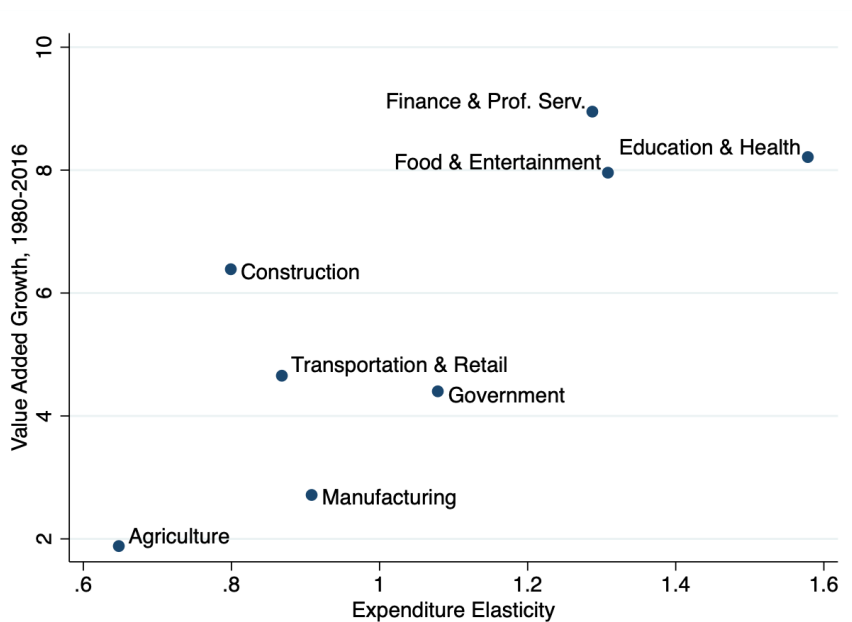
¹See the discussion of the related literature at the end of this section for the evidence on labor-market polarization.

²As we further discuss in the related literature subsection, our mechanism is related but distinct from the spillover hypothesis in Manning (2004), Mazzolari and Ragusa (2013), and Leonardi (2015).

³We use the estimation method developed in Aguiar and Bils (2015), with expenditures defined over value-added consumption as in Buera et al. (2015).

⁴We follow Acemoglu and Autor (2011) and classify occupations into three skill categories (high-, middle- and low-skill) that are defined according to their average wage in 1980. See Appendix B for details. These correlation patterns also hold if we measure factor intensity with employment shares or if we compute the correlation weighted by the sector value-added share in total value added. We prefer this measure of factor intensity because it has a direct link with the model presented in Section 2.

Figure 1: Sectoral Nominal Value-Added Growth 1980-2016 and Expenditure Elasticity



Notes: Sectoral value added measures are from the BEA. Expenditure elasticities are estimated from cross-sectional household data using the [Aguiar and Bils \(2015\)](#) methodology (see Section 3.1).

over the 1980-2016 period.⁵ This implies that the correlation between expenditure elasticities and factor intensities across sectors is also remarkably stable over this period.⁶

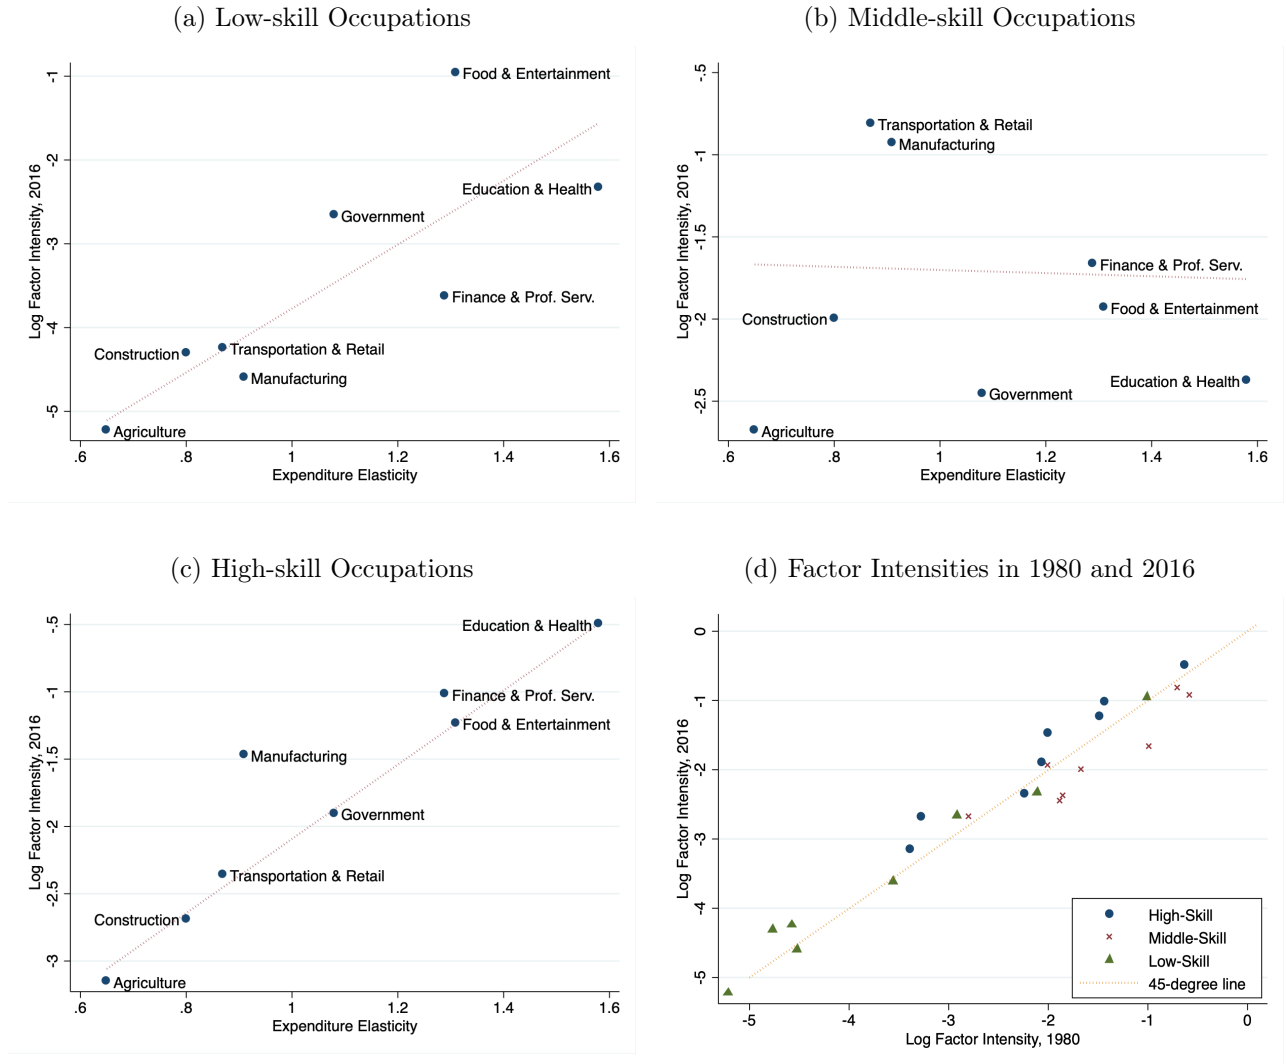
We develop and quantify a multi-sector general equilibrium model to assess the importance of the income-driven channel for the polarization of labor-market outcomes. We gradually introduce the elements of the model to better understand their importance for our mechanism. In Section 2, we present the sectoral production functions, which allow us to define factor intensities. We then conduct a shift-share decomposition of the wage-bill shares across occupations. This exercise reveals that changes in the sectoral composition of the economy (captured by the share term) lead to sizable reallocation in occupational wage-bill shares. They account for 73%, 28% and 31% of the changes in the wage-bill shares of low-, middle- and high-skill occupations from 1980 through 2016. In addition, we find that the share terms account for 78%, 40% and 30% of the changes in the employment shares of low-, middle-, and high-skill occupations over this span.

Changes in the sectoral composition of the economy can be driven by many factors, including some that may be endogenous to polarization such as technological change or off-

⁵Figure 2d also shows how factor intensities of high-skill occupations tend to be above the 45-degree line, while for middle-skill occupations they tend to be below. This is consistent with biased technological change and offshoring affecting these occupations. In contrast, low-skill occupations tend to be closer to the 45-degree line, suggesting that these mechanisms may have been less important for the low-skilled.

⁶See Table A.1 in Appendix A.

Figure 2: Factor Intensity and Expenditure Elasticities across Sectors



Notes: The factor intensity of occupation j in sector s is the wage bill of occupation j in sector s over total value added in sector s . Employment is computed using hours worked from the Census/ACS, wages are from the CPS and value added, from the BEA (see Appendix B). Expenditure elasticities are estimated from household data using the Aguiar and Bils (2015) methodology (see Section 3.1).

shoring. To isolate the income-driven channel, we introduce in Section 3.1 a nonhomothetic demand system with constant elasticity of substitution across sectors. We estimate the demand elasticities using household-level data from the Consumer Expenditure Survey (CEX) supplemented with imputations of household consumption of public health, education and finance charges not covered by the CEX. Importantly, by disciplining the demand side of our model with micro-level estimates, we avoid using data on aggregate moments of interest to calibrate critical preference parameters. Before turning to our general-equilibrium analysis, we use the estimated demand in a partial-equilibrium setting to quantify the role of

income and price effects in generating changes in the sectoral composition of the economy and labor-market polarization. Section 3.3 reports the results of feeding the observed changes in expenditure and prices to our estimated demand. We find that increases in expenditure account for the bulk of changes in the sectoral composition of the economy and the share component of the shift-share decomposition of occupational wage bills.

Section 4 presents our full-fledged quantification of the income-driven polarization channel. The model embeds the preferences and production functions introduced in the previous sections into a general-equilibrium setting, so that prices and income are endogenous. The model also incorporates an occupational choice by workers that allows us to study distributional outcomes in terms of both hours worked and relative wages. Traditional explanations of polarization are reflected in time-varying, occupation-sector-specific factor intensities and sector-specific total factor productivity (TFP). We calibrate these parameters to match the change in value-added shares from 1980 to 2016. The income-driven polarization channel results from exogenous increases in neutral TFP and average worker productivity across occupations. These changes are calibrated to match the increase in aggregate nominal expenditures per capita and the change in the personal consumption expenditure deflator. We find that the income-driven channel explains 90% of the observed change in the wage-bill share of low-skill occupations and 35% of the change in this share of high-skill occupations from 1980 through 2016; however it only accounts for 2% of middle-skill occupations. This contribution to the polarization of the wage-bill distribution reflects polarization in both hours worked and relative wages. Our calibrated model predicts that the income-driven channel accounts for 64%, 35% and 28% of the change in the share of hours worked by low-, middle-, and high-skill occupations, as well as for 46% and 29% of the increase in the relative wage of low- and high- to middle-skill occupations.

Taken together, our results suggest that a significant part of the polarization of labor markets observed during 1980-2016 in the US is due to the income-driven channel. We find that the income-driven channel is the main driver for low-skill occupations. In contrast, the evolution of sectoral technologies of production account for the vast majority of the changes for middle-skill occupations. For high-skill occupations, we find that the income-driven channel accounts for around one-third of the overall changes. Our findings are robust to a number of extensions, such as, allowing for a non-unitary elasticity of substitution between occupations in production; accounting for the wedges between sectoral production and expenditures introduced by international trade; and shifting the rules used to assign capital income across households.

After having established the importance of the income-driven channel for the US since the 1980s, we explore in Section 5 the relevance of the mechanism in other geographies and time

periods. Labor-market polarization has been documented in Western Europe (Goos et al., 2009, 2014). We use our model to study the role of nonhomotheticities in the evolution of labor-market polarization in Western European economies since 1980. We find that the income-driven channel contributes to polarization in these countries by a similar magnitude to what we found in our baseline results for the US. We conclude our paper by looking again to the US and exploring the relevance of our mechanism for the 1950-1980 period (Bárány and Siegel, 2018, suggest an earlier onset of polarization). We also use our model to project the effect of the income-driven channel on the transformation of labor markets from 2016 to 2035. Our model suggests that the income-driven polarization of labor markets will continue at least at the same pace as we have observed during the last 25 years.

Related literature Our paper is related to two different strands of the literature, one studying labor-market polarization and the other, structural change and inequality. First, our work is related to a vast literature in labor economics that has documented and explored the drivers of labor-market polarization as initially suggested by Acemoglu (1999). From the work by Autor et al. (2003), Goos and Manning (2007), Goos et al. (2009, 2014), Autor et al. (2006, 2008), Autor and Dorn (2009, 2013) and others, technical change has emerged as the leading explanation for labor-market polarization, with offshoring playing a secondary role.⁷ Goos and Manning (2007) built on the task framework developed in Autor et al. (2003) and proposed the “routinization” hypothesis, whereby computer capital substituted for workers in routine-intensive occupations—which tend to be middle-skill occupations. Our paper is most closely related to Autor and Dorn (2013), Goos et al. (2014) and, more recently, Cortes et al. (2017). These papers have incorporated the routinization hypothesis into a modelling framework. Like these papers, we model goods demanded by consumers as gross substitutes (generating a Baumol’s cost-disease mechanism) and sectors differ in terms of their factor intensity and their exposure to technological progress. The key difference is that our model introduces the income-driven channel (through nonhomotheticities in preferences).⁸

This strand of the literature has deemed the role of income effects to be small in accounting for labor-market polarization (Manning, 2004, Mazzolari and Ragusa, 2013, and Leonardi, 2015). These papers emphasize a spillover or trickle-down effect, in which the demand for low-skill services increases as a result of rich households becoming richer.⁹ Their empirical

⁷Basco and Mestieri (2013) document that trade costs have declined more in middle-skill intensive industries (which are also more intensive in information technologies and routine tasks). They show that this can lead to wage polarization in rich countries. See the references therein for trade-based accounts of polarization. See Cozzi and Impullitti (2016) for a dynamic account of polarization and trade in a growth model.

⁸Like Goos et al. (2014), we structure our results and model quantification around shift-share decompositions and emphasize that both shift and share components play an important role in accounting for polarization.

⁹Leonardi (2015) uses household expenditure data to show that educated households consume more high-skill and, to a lesser extent, very low-skill services. Leonardi then links an “education-specific elasticity of

strategy, as well as the analysis of income effects in [Autor and Dorn \(2013\)](#), typically identifies income effects by comparing the increase in demand for low-skill services between commuting zones in which top incomes have increased substantially and those in which top incomes have increased more modestly. These studies find little support for this mechanism. While this trickle-down theory is consistent with our mechanism, it misses the *level* effect that rising average household income has on the sectoral composition of demand. That is, a difference-in-differences estimation cannot identify a common effect across commuting zones driven by an increase in the average household income. In contrast, we use cross-sectional household variation to estimate expenditure elasticities for different sectors.¹⁰ We then use our estimated demand system to feed in the observed increases in household expenditures and prices. We find that the income-driven channel is significant—especially for low-skill occupations.

Our paper is also related to a small (but growing) complementary literature in macroeconomics that has linked structural change and inequality. To our knowledge, [Lee and Shin \(2017\)](#) and [Bárány and Siegel \(2018, 2019\)](#) are the few papers that have studied the role of the sectoral composition of the economy for labor market polarization.¹¹ In contrast to our paper, these papers assume that structural change is solely driven by technological progress; there is no income-driven mechanism like ours in this previous research. For this reason, our paper finds that structural change has played a more substantial role in accounting for labor-market polarization than these papers. Importantly, the model in our paper includes a conceptually different mechanism, which yields different testable predictions and is absent in this previous research. For example, to the extent that household income keeps rising, our framework identifies a novel economic force that is likely to drive polarization in the future.

Other papers in the literature on structural change and inequality have focused on the skill premium, starting with [Schimmelpfennig \(1998\)](#) and, more recently, [Buera and Kaboski \(2012\)](#), [Cravino and Sotelo \(2017\)](#), [Buera et al. \(2018\)](#), [Caron et al. \(2020\)](#) and [Ngai and Sevinc \(2020\)](#).¹² The paper that is closest to ours is [Buera et al. \(2018\)](#). They calibrate a two-skill, two-sector model and infer the contribution of structural change to the rise in the skill premium from the calibrated growth in sectoral TFPs. Instead, using a three-skill,

demand” to the rise of the skill premium in a static general equilibrium model. He finds a small effect for demand shifters in his calibrated model.

¹⁰In the appendix of [Goos et al. \(2014\)](#), the authors augment a CES demand estimation with income and population. In general, they find the effects of income and population to be statistically insignificant. In contrast to their exercise, our demand specification is derived from a demand system, uses household variation rather than aggregate time series, and is identified through an instrumental-variable strategy.

¹¹[Cerina et al. \(2017\)](#) link polarization in the US to increases in female labor force participation.

¹²Using a shift-share design in German data, [Schimmelpfennig](#) shows that structural change accounts for 40% of the rise in the skill premium once the input-output structure of the economy is taken into account. [Cravino and Sotelo \(2017\)](#) and [Caron et al. \(2020\)](#) study the effect of global growth and international trade on structural change and the evolution of the skill premium.

eight-sector model, we quantify the contribution of the income-driven channel to polarization in wage bills, employment, and relative wages, isolating this mechanism from the effect of relative TFP growth across sectors. Finally, our modelling of nonhomotheticities focuses on the overall expenditure elasticity. An alternative formulation of nonhomotheticity relates income to the quality of goods consumed within a sector, as recently emphasized in [Jaimovich et al. \(2019\)](#). This alternative manifestation of nonhomotheticities may further enhance the effect of the income-driven channel on the polarization of labor markets.

2 Production and Wage-Bill Decompositions

This section presents the production side of our multi-sector model. We show that the production side of the model suffices to quantify the contribution of sectoral reallocation to wage-bill polarization. We use this insight to present a theory-grounded shift-share decomposition.

2.1 Production

We consider an economy with $\mathcal{S} = \{1, \dots, S\}$ sectors. Each sector $s \in \mathcal{S}$ produces output Y_{st} at time t according to the value-added constant-returns-to-scale Cobb-Douglas production function

$$Y_{st} = A_{st} \left(\prod_{j \in \{L, M, H\}} X_{jst}^{\alpha_{jst}} \right)^{\beta_{st}} K_{st}^{1-\beta_{st}}, \quad (1)$$

where $\alpha_{st}, \beta_{st} \in (0, 1)$, $\sum_{j \in \{L, M, H\}} \alpha_{jst} = 1$, K_{st} is the capital stock, and X_{jst} is the number of hours from occupation j used in sector s . We consider three occupations: *Low*-, *Middle*- and *High*-skill. A_{st} is a Solow residual that captures both Hicks-neutral technological improvements and capital-skill complementarities not explicitly modeled. We also allow the parameters $\{\alpha_{jst}, \beta_{st}\}_{s \in \mathcal{S}, j \in \{L, M, H\}}$ to vary over time to parsimoniously capture the patterns of substitution across factors of production implied by biased technological change, routinization, offshoring, changes in production wedges, etc.^{13,14}

¹³To illustrate this point, consider a general constant-returns-to-scale production function $Y_{st} = F_s(\{X_{jst}\}, K_{st})$. Given a price of sectoral output p_{st} and wages w_{jt} , the firm's first order conditions implied by profit maximization are $w_{jst}X_{jst} = \hat{\alpha}_{jst}p_{st}Y_{st}$ where $\hat{\alpha}_{jst} = \partial F_s / \partial X_{jst} \cdot X_{jst} / Y_{st}$. Therefore, the factor intensity $\hat{\alpha}_{jst}$ is a well-defined object that can be measured by the wage bill of occupation j in sector s at time t . For example, consider the production function from Section 4.4, where we have a CES aggregator across occupations, $Y_{st} = A_{st} \left(\sum_{j \in \{L, M, H\}} \gamma_{jst}^{\frac{1}{\sigma}} X_{jst}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\beta_{st}\sigma}{\sigma-1}} K_{st}^{1-\beta_{st}}$. The wage bill paid to occupation j is $w_{jst}X_{jst} = \beta_{st} \left(\gamma_{jst} \left(\frac{w_{jt}}{\sum_j \gamma_{jst} w_{jt}^{1-\sigma}} \right)^{1-\sigma} \right) p_{st}Y_{st}$, where $\alpha_{jst} \equiv \gamma_{jst} \left(\frac{w_{jt}}{\sum_j \gamma_{jst} w_{jt}^{1-\sigma}} \right)^{1-\sigma}$ and $\hat{\alpha}_{jst} = \alpha_{jst}\beta_{st}$.

¹⁴Capital-skill complementarity has been emphasized, among others, by [Caselli \(1999\)](#) and [Krusell et al. \(2000\)](#). [Lee and Shin \(2017\)](#) and [Bárány and Siegel \(2019\)](#) have emphasized the importance of sector-

Under perfect competition, given a price of sectoral output p_{st} and wages w_{jt} , the demand for occupation j in sector s at time t is

$$WB_{jst} \equiv w_{jt}X_{jst} = \hat{\alpha}_{jst}VA_{st}, \quad (2)$$

where $\hat{\alpha}_{jst} \equiv \alpha_{jst}\beta_{st}$, WB_{jst} denotes the wage bill accrued by occupation j in sector s at time t , and VA_{st} denotes the nominal value added in sector s (i.e., $VA_{st} = p_{st}Y_{st}$). It follows from Equation (2) that $\hat{\alpha}_{jst}$ captures the intensity of occupation j in sector s at time t ,

$$\frac{WB_{jst}}{VA_{st}} = \hat{\alpha}_{jst}. \quad (3)$$

This is the measure of factor intensity that we have presented in Figure 2 of the Introduction. The wage-bill share of occupation j in sector s can also be expressed in terms of the factor intensity $\hat{\alpha}_{jst}$ as

$$\frac{WB_{jst}}{WB_{st}} = \frac{\hat{\alpha}_{jst}}{\beta_{st}}, \quad (4)$$

where β_{st} is the labor share in sector s value added. Expression (4) shows that the wage-bill share of occupation j in sector s only depends on factor intensities. We use this property extensively in what follows. Next, denoting total employment in occupation j by $X_{jt} \equiv \sum_{s \in \mathcal{S}} X_{jst}$, the aggregate wage bill of occupation j can be obtained summing (2) across sectors,

$$WB_{jt} \equiv w_{jt}X_{jt} = \sum_{s \in \mathcal{S}} \hat{\alpha}_{jst}VA_{st}. \quad (5)$$

As we show next, these optimality conditions on the production side of the economy suffice to assess the importance of sectoral reallocation on wage-bill polarization.

2.2 Decomposition of the Evolution of Wage Bills

Using Equation (5), we can express the total wage bill of occupation j relative to the total wage bill in the economy as

$$\frac{WB_{jt}}{WB_t} = \sum_{s \in \mathcal{S}} \hat{\alpha}_{jst} \frac{VA_{st}}{VA_t} \frac{VA_t}{WB_t} = \sum_{s \in \mathcal{S}} \frac{\hat{\alpha}_{jst}}{\beta_t} \frac{VA_{st}}{VA_t}, \quad (6)$$

where in the last equality we have denoted the labor share of the economy by $\beta_t \equiv WB_t/VA_t$. Equation (6) shows that the aggregate wage-bill share of occupation j depends both on the

occupation specific technological progress.

sectoral composition of the economy, VA_{st}/VA_t , and factor intensities, $\hat{\alpha}_{st}/\beta_t$. In particular, if sectors that are more intensive in occupation j grow faster than the average, the aggregate wage-bill share of this occupation will increase. Indeed, this is what Figures 1 and 2 in the Introduction suggest for low- and high-skill occupations.

This finding contrasts with the *sectoral* wage-bill share of occupation j derived in Equation (4), which only depends on factor intensities. This comparison is relevant because in a one-sector economy, $S = 1$, the aggregate wage-bill share of occupation j is also given by Equation (4). Hence, in a one-sector setting, changes in the distribution of the wage bill across occupations can only come from changes in factor intensities $\{\hat{\alpha}_{jst}\}$. However, as originally pointed out by Schimmelpfennig (1998) and discussed above, recognizing that the economy is populated by multiple sectors that may grow at different rates opens the possibility for changes in the sectoral composition of the economy to affect the distribution of wage bills across occupations.

To explore this possibility, we express the change in the wage-bill share of occupation j from time 0 to t as

$$\frac{WB_{jt}}{WB_t} - \frac{WB_{j0}}{WB_0} = \sum_{s \in \mathcal{S}} \frac{\hat{\alpha}_{jst}}{\beta_t} \frac{VA_{st}}{VA_t} - \sum_{s \in \mathcal{S}} \frac{\hat{\alpha}_{js0}}{\beta_0} \frac{VA_{s0}}{VA_0}. \quad (7)$$

Using a standard shift-share approach,¹⁵ we can decompose the change in the wage-bill share of occupation j as

$$\frac{WB_{jt}}{WB_t} - \frac{WB_{j0}}{WB_0} = \underbrace{\sum_{s \in \mathcal{S}} \left(\frac{\hat{\alpha}_{jst}}{\beta_t} - \frac{\hat{\alpha}_{js0}}{\beta_0} \right) \frac{VA_{s0}}{VA_0}}_{Shift} + \underbrace{\sum_{s \in \mathcal{S}} \frac{\hat{\alpha}_{jst}}{\beta_t} \left(\frac{VA_{st}}{VA_t} - \frac{VA_{s0}}{VA_0} \right)}_{Share}. \quad (8)$$

The first summation in (8) is the shift term. It captures the effect of changes in factor intensities of occupation j on the wage-bill share, holding the sectoral distribution of value added constant. The second summation corresponds to the share term. It captures the effect of the change in the sectoral composition of value added on the change in the wage-bill share for occupation j , holding factor intensities constant.¹⁶

¹⁵We thank Richard Rogerson for suggesting using a shift-share decomposition. Alternative decompositions, e.g., based on wage-bill growth yield very similar conclusions.

¹⁶This is our preferred decomposition of the wage bill, since it is grounded in theory. We find that using the reduced-form decomposition of the wage bill

$$\frac{WB_{jt}}{WB_t} - \frac{WB_{j0}}{WB_0} = \sum_s (\alpha_{jst} - \alpha_{js0}) \frac{WB_{s0}}{WB_0} + \sum_s \alpha_{jst} \left(\frac{WB_{st}}{WB_t} - \frac{WB_{s0}}{WB_0} \right) \quad (9)$$

yields quantitatively similar results. The key difference with the decomposition in (8) is that, in this reduced-form decomposition, changes in aggregate and sectoral labor shares affect the share term.

2.3 Assessing the Contribution of Sectoral Composition

To quantify the contribution of each term of the shift-share decomposition in Equation (8), we group the US economy into eight broad sectors.¹⁷ We take the nominal sectoral value added from the BEA. The occupation intensities for each sector, $\hat{\alpha}_{jst}$, are computed from Equation (2) combining information on the share in the sectoral wage bill of each occupation, α_{jst} , with sectoral labor shares, β_{st} .¹⁸ Motivated by the US studies of polarization (e.g., Autor et al., 2003), we start our analysis in 1980 and take the latest date with data availability, 2016, as our final year.

We observe a stark polarization in wage bills over the 1980-2016 period. The shares of high- and low-skill occupations in the total wage bill increased by 19 and 2 percentage points. In contrast, the share of middle-skill occupations declined by 21 percentage points. Table 1 reports the shift-share decomposition derived in Equation (8) of these changes in wage-bill shares. The share term, which captures the role of changes in the sectoral composition of value added, accounts for a significant fraction of the observed changes in occupational wage-bill shares: 31% of the increase in the wage-bill share of high-skill occupations, 28% of the decrease in the wage-bill share of middle-skill occupations, and 73% of the increase in the wage-bill share of low-skill occupations.¹⁹ Changes in the factor intensities of occupations and labor account, through the shift term, for the complementary fractions. In sum, this shift-share decomposition shows that changes in the sectoral composition of the economy play a significant role for the change in the wage-bill shares of high- and middle-skill occupations and a dominant role for the change in the wage-bill share of low-skill occupations.

¹⁷The eight sectors are i) agriculture, ii) manufacturing, mining and utilities, iii) construction and real estate iv) retail and wholesale trade, and transportation, v) finance, insurance, information, professional and other services, vi) health and education, vii) food, arts and entertainment, and viii) government. This grouping is constructed from the 15 sectors in the BEA's input-output tables after considering both the traditional aggregation of sectors and the estimates of the income elasticity of demand for these 15 sectors. See Appendix B for more details.

¹⁸Sectoral wage-bill shares of each occupation are computed using wage data for each occupation from the CPS and hours worked for each occupation in each sector from the ACS, as in Acemoglu and Autor (2011). We compute sectoral labor shares as

$$\beta_{st} = \left(\frac{WB_{st}^{CPS/ACS}}{WB_t^{CPS/ACS}} \right) \left(\frac{VA_t^{BEA}}{VA_{st}^{BEA}} \right) \left(\frac{WB_t^{BEA}}{VA_t^{BEA}} \right), \quad (10)$$

where $\frac{WB_{st}^{CPS/ACS}}{WB_t^{CPS/ACS}}$ is the share of sector s wage bill in the total wage bill when these are computed using the hours worked from the ACS and the wages per hour from the CPS, $\frac{VA_t^{BEA}}{VA_{st}^{BEA}}$ is the inverse of the share in value added of sector s from the BEA; $\frac{WB_t^{BEA}}{VA_t^{BEA}}$ is the aggregate labor share from the BEA. This procedure allows us to ensure that we match the aggregate BEA labor share.

¹⁹The results from the reduced-form decomposition of the wage-bill shares described in Equation (9) are quite similar. The share term accounts for 29% and 88% of the increase in the wage-bill share for high- and low-skill occupations, as well as for 35% of the decline in the wage-bill share of middle-skill occupations.

Table 1: Shift-Share Decomposition of US Wage Bill and Employment, 1980-2016

	Wage Bill Share			Employment Sh.		
	Low	Middle	High	Low	Middle	High
Total Change	0.021	-0.21	0.19	0.034	-0.16	0.13
Shift	27%	72%	69%	22%	60%	70%
Share	73%	28%	31%	78%	40%	30%

The importance of the share term is due to the substantial changes in the composition of the US economy over the 1980-2016 period (documented in Figure 1), as well as the correlations between the change in sectoral value-added shares and occupation intensities. The correlations between the change in sectoral value-added shares and the occupation intensity of high-, middle- and low-skill occupations are 0.40, -0.52 and 0.15. These correlation patterns imply that, as the sectoral composition of the economy changes, the relative demand for high- and low-skill occupations increases, while the relative demand for middle-skill occupations declines. These changes in the relative demand for each occupation are reflected in the changes in their wage-bill shares.

Decomposing employment shares Even though the production side of our model is not sufficient on its own to derive predictions for the distribution of hours worked, we can still use a reduced-form shift-share decomposition of hours worked across occupations to explore the relevance of the sectoral changes in the economy. Let the number of hours worked in sector s be $X_{st} \equiv \sum_{j \in \{L, M, H\}} X_{jst}$, and $X_t \equiv \sum_{s \in \mathcal{S}} X_{st}$, the total number of hours worked in the economy. We can decompose the change in the share of hours worked by occupation j as

$$\frac{X_{jt}}{X_t} - \frac{X_{j0}}{X_0} = \underbrace{\sum_{s \in \mathcal{S}} (\rho_{jst} - \rho_{js0}) \frac{X_{s0}}{X_0}}_{Shift} + \underbrace{\sum_{s \in \mathcal{S}} \rho_{jst} \left(\frac{X_{st}}{X_t} - \frac{X_{s0}}{X_0} \right)}_{Share}, \quad (11)$$

where $\rho_{jst} \equiv \frac{X_{jst}}{X_{st}}$ is the share of sector s total hours corresponding to occupation j .

The first row of results in the second panel in Table 1 shows that the distribution of hours worked across occupations has experienced a polarization of similar magnitude to that of the wage bill. The shares of hours corresponding to high- and low-skill occupations have increased by 13 and 3 percentage points, while the share of hours of middle-skill occupations has declined by 16 percentage points. The shift-share decomposition suggests that changes in the sectoral composition of employment play a significant role in high- and middle-skill occupations, as well as a dominant role in low-skill occupations. Specifically, the share term accounts for 30% of the increase in the share of hours worked in high-skill occupations, 40% of

Table 2: Shift-Share Decomposition for Alternative Aggregations, US 1980-2016

	Wage Bill Share			Employment Sh.		
	Low	Middle	High	Low	Middle	High
Total Change	0.021	-0.21	0.19	0.034	-0.16	0.13
Share/Total Change (%) with:						
8 Sectors, 3 Occupations	73%	28%	31%	78%	40%	30%
2 Sectors, 3 Occupations	49%	3%	19%	46%	23%	17%
2 Sectors, 2 Occupations		-3%	19%		17%	17%

the decline in middle-skill occupations, and 78% of the increase in low-skill occupations. We conclude from this exercise that changes in the sectoral composition of the economy contribute similarly to the observed changes in hours worked and wage bills across occupations.

Sectoral and Occupational Disaggregation We finish this section by exploring the importance of using disaggregated sectors and occupations in our setting. Table 2 presents the contribution of the share term across alternative, coarser aggregations of sectors and skills. The first and second rows report the observed changes in wage bill and employment shares, and the contribution of the share term in our baseline classification of eight sectors and three occupations. The third row contains the contribution of the share term when we aggregate our eight sectors into service- and goods-producing sectors.²⁰ The fourth row maintains the two-sector classification and combines the middle- and low-skill occupations into a single group of unskilled occupations. The main finding is that aggregating up sectors and occupations greatly diminishes the contribution of the share term—it even reverses the sign of the contribution to the wage bill in the case with only two skill levels of occupations. We conclude from this exercise that it is quantitatively important to implement a disaggregated classification of sectors and occupations to properly capture the actual contribution of the sectoral composition of the economy to labor-market polarization.

3 Preferences and Partial-Equilibrium Quantification

This section introduces the demand side of our model. Our setting allows for supply side forces (e.g., biased technical change, robotization, offshoring) that affect the sectoral composition of demand through changes in sectoral relative prices, and it also allows for income effects that affect the composition through nonhomotheticities in demand. After presenting

²⁰In the service-producing group, we include health and education, retail and wholesale trade, finance, insurance, real estate, information and professional services, food and entertainment, and government. In the goods-producing group we include agriculture, mining and utilities, manufacturing, and construction.

the estimation of preference parameters, we use the restrictions imposed by the partial equilibrium of the model to disentangle the contributions of income and price effects to changes in the sectoral distribution of value added and wage-bill shares.

In our exposition, we make two simplifying assumptions. First, we assume that all income accrued by the workers is spent in consumption, abstracting from savings. Second, we study an economy with a representative household. That is, all the resources earned by the household are pulled together and, given the prevailing prices, the household decides its consumption bundle. We relax this assumption in Section 4, where we develop the general equilibrium version of the model and introduce heterogeneity in consumption across households.

3.1 Nonhomothetic Constant Elasticity of Substitution Preferences

The utility of the representative household at time t , U_t , is a nonhomothetic Constant Elasticity of Substitution (CES) aggregator defined over consumption goods $\{c_{st}\}_{s \in \mathcal{S}}$ through the constraint

$$\sum_{s \in \mathcal{S}} (U_t^{\varepsilon_s} \zeta_s)^{\frac{1}{\sigma}} c_{st}^{\frac{\sigma-1}{\sigma}} = 1. \quad (12)$$

The parameter σ is the (constant) elasticity of substitution across goods, the nonhomotheticity parameter ε_s controls the expenditure elasticity of sector s , and ζ_s captures the constant taste component over s . We focus on the empirically relevant case of gross complements, $0 < \sigma < 1$, and $\zeta_s, \varepsilon_s > 0$ for all $s \in \mathcal{S}$.²¹ Intuitively, this formulation of preferences allows for the level of utility $U_t^{\varepsilon_s}$ to enter asymmetrically in (12) as an additional taste component. In contrast to ζ_s , this term is variable and endogenously determined. As a result, the overall weight attached to the consumption of good s , $U_t^{\varepsilon_s} \zeta_s$, depends on the level of utility itself, U_t , with an elasticity controlled by ε_s . If ε_s 's were constant across all s , we would obtain homothetic CES preferences.

Given a set of prices $\{p_{st}\}_{s \in \mathcal{S}}$ and total expenditure E_t , a household maximizing utility (12) subject to the budget constraint $\sum_{s \in \mathcal{S}} p_{st} c_{st} \leq E_t$ chooses $\{c_{st}\}_{s \in \mathcal{S}}$ so that

$$c_{st} = \zeta_s \left(\frac{E_t}{p_{st}} \right)^{\sigma} U_t^{\varepsilon_s}. \quad (13)$$

The corresponding expenditure function is given by $E_t^{1-\sigma} = \sum_{s \in \mathcal{S}} \zeta_s U_t^{\varepsilon_s} p_{st}^{1-\sigma}$.

We can normalize one taste parameter $\zeta_s \equiv 1$ (as with homothetic CES) and one income

²¹A sufficient condition for these preferences to be well-defined is that $\varepsilon_s > 0$ if $0 < \sigma < 1$, and $\varepsilon_s < 0$ if $\sigma > 1$ for all $s \in \mathcal{S}$. See Hanoch (1975) and Comin et al. (2015) for further discussion.

elasticity parameter $\varepsilon_s \equiv 1$ for some $s \in \mathcal{S}$ (see Comin et al., 2015). These normalizations cardinalize (12) and uniquely define a cost-of-living index P_t and a real consumption index C_t of the representative household, $U_t = \frac{E_t}{P_t} \equiv C_t$. The cost-of-living index can be expressed in terms of observables and demand parameters as

$$P_t = \left[\sum_{s \in \mathcal{S}} (\zeta_s p_{st}^{1-\sigma})^{\theta_s} (x_{st} E_t^{1-\sigma})^{1-\theta_s} \right]^{\frac{1}{1-\sigma}}, \quad (14)$$

where $x_{st} = p_{st} c_{st} / E_t$ denotes the expenditure share in sector s , and $\theta_s \equiv (1 - \sigma) / \varepsilon_s$. With this notation, the expenditure share in sector s is

$$x_{st} = \zeta_s \left(\frac{p_{st}}{P_t} \right)^{1-\sigma} C_t^{\varepsilon_s - (1-\sigma)}. \quad (15)$$

Finally, the expenditure elasticity of sector s is

$$\eta_{st} \equiv \frac{\partial \ln p_{st} c_{st}}{\partial \ln E_t} = \sigma + (1 - \sigma) \frac{\varepsilon_s}{\sum_{s \in \mathcal{S}} x_{st} \varepsilon_s}. \quad (16)$$

Thus, whether a good has an expenditure elasticity higher (or lower) than 1 depends on whether $\varepsilon_s > (<) \sum_{s \in \mathcal{S}} x_{st} \varepsilon_s$, which depends on the total level of expenditure of the household, E_t . This implies that the same good can be a luxury or a necessity depending on the level of expenditure of a household.

Equilibrium and Channels of Sectoral Reallocation We close the model by imposing the market clearing condition that, in each sector, the value of consumption equals production. Using the demand for good s , Equation (13), market clearing implies that nominal value added in sector s is

$$VA_{st} = p_{st} c_{st} = \zeta_s p_{st}^{1-\sigma} E_t^{\sigma + \varepsilon_s} P_t^{-\varepsilon_s}, \quad (17)$$

where the total expenditure of the representative household, E_t , is equal to total labor income plus capital income (net of depreciation),

$$E_t = \sum_{s \in \mathcal{S}} \sum_{j \in \{L, M, H\}} w_{jt} X_{jst} + r_t K_t. \quad (18)$$

Equation (17) illustrates that, in our model, the evolution of sectoral value added is driven by two forces: changes in aggregate expenditures, E_t , and changes in sectoral prices, $\{p_{st}\}$.²²

²²Note from Equation (14) that P_t is itself a function of aggregate expenditure and prices.

Supply-side drivers of polarization and inequality such as biased technical change, automation, de-unionization or offshoring affect the sectoral composition of value added through their impact in relative sectoral prices.²³ Indeed, in a model with homothetic preferences, changes in relative prices are the only source of sectoral reallocation. To see this, we use Equation (17) and consider the ratio of demand for two sectors, s and s' , when preferences are homothetic (i.e., $\varepsilon_s = 1$ for all $s \in \mathcal{S}$),

$$\frac{VA_{st}}{VA_{s't}} = \frac{\zeta_s}{\zeta_{s'}} \left(\frac{p_{st}}{p_{s't}} \right)^{1-\sigma}.$$

In this case, changes in the relative sectoral composition depend on the evolution of relative prices and are independent of the overall level of expenditure, E_t .

Nonhomotheticities introduce a distinct, income-driven mechanism that affects the evolution of the sectoral composition of value added. As expenditure increases, consumers shift the composition of expenditure from low expenditure elastic (low ε_s) to high expenditure elastic sectors (high ε_s). Even if relative prices were equal across sectors, our demand system would imply changes in the sectoral composition of the economy driven by aggregate expenditure. To see this, using again Equation (17), we have that the relative sectoral demand when $p_{st} = p_{s't}$ becomes

$$\frac{VA_{st}}{VA_{s't}} = \frac{\zeta_s}{\zeta_{s'}} \left(\frac{E_t}{P_t} \right)^{\varepsilon_s - \varepsilon_{s'}}.$$

Both in partial equilibrium (Section 3.3) and in general equilibrium (Section 4), the impact of the income-driven channel on polarization operates through the effects of expenditure on the sectoral composition of the economy. Before conducting any quantification of this channel, we need to estimate the preference elasticity parameters that drive the income and price effects, $\{\varepsilon_s\}_{s \in \mathcal{S}}$ and σ .

3.2 Household-level Data and Estimation of Demand Elasticities

We use household-level data to provide evidence on heterogeneity in expenditure elasticities across broad sectors of the economy and estimate the elasticities governing our demand system. This latter exercise will serve as a basis for the quantification of the model. Since our estimation borrows extensively from previous literature (Aguar and Bils, 2015 and Comin et al., 2015 in particular), we provide a relatively concise description of our data and estima-

²³This account captures the effect of offshoring and international trade on labor demand through sectoral prices and factor intensities. Additionally, offshoring and trade can affect sectoral demand by breaking the market-clearing condition that domestic production equals domestic consumption. We study this channel in Section 5.

tion exercises in the main text and relegate the details to Appendix B.2.

Data Description Our main data source is the Consumption Expenditure survey (CEX) for years 2000 and 2001. The CEX is a rotating panel that contains detailed information on household quarterly expenditures and characteristics. We follow Aguiar and Bils (2015) in their sample selection, which is standard and, in turn, very similar to the previous literature (e.g., Krueger and Perri, 2006, among others). For example, we restrict our attention to urban households with a present household head. We only depart from the standard sample selection in that we do not restrict the age of the reference household, thus including all households older than 64. We do this to account for potentially important healthcare expenditures in the late stages of life.

Since our model is specified over value-added consumption, we need to transform the reported household expenditures into value-added consumption, as in Herrendorf et al. (2013) and Buera et al. (2015, 2018). We follow the procedure described in Buera et al. (2015) and map each of the over nine-hundred CEX expenditure categories to the 76 lines appearing in the NIPA Personal Consumption Expenditures table. From here, we can map these categories to the industries appearing in the 2000 BEA Input-output table and infer the value-added consumption of each industry embodied in the CEX expenditures. In short, this exercise breaks up expenditures into its value-added components. For example, expenditure in restaurants is split in its value-added components coming from agriculture (food), manufacturing (processing of the food), food services (cooks and waiters), utilities (electricity usage), etc.

One limitation of the CEX data is that it only measures out-of-pocket expenditures. Thus, it likely provides an underestimate of expenditure categories that are partially publicly provided, such as education and health services. Since we will use cross-sectional household variation to estimate our demand elasticities, we could obtain biased estimates if out-of-pocket expenditures for a given expenditure category were correlated with household total expenditure. To partially address these concerns, we impute public expenditures in K-12, Medicare and Medicaid to our household data. We use information on average expenditure per pupil by school district and expenditure per patient by hospital referral region and match it to households based on the number of school-age children and number of recipients of Medicare and Medicaid. We also impute fund management expenses based on reported households' pension funds and total value of stocks, bonds and other securities.²⁴ We take

²⁴For public primary and secondary education, we take expenditure per pupil at the school district level from the Common Core Database and regress it on household county median income, state fixed effects and a time trend. This exercise gives a predicted expenditure per pupil as a function of household income, location and time. We then use information on the number of kids in K-12 age of a household and impute consumption accordingly for those households that report zero expenditures in elementary and high-school tuition. For Medicare and Medicaid data, we use average expenditure per patient by hospital referral region from the

this augmented version of our data as our baseline, since it features the most comprehensive account of household expenditures.²⁵

Reduced-form Evidence on Expenditure Elasticities We begin our analysis by providing reduced-form evidence on heterogeneity in expenditure elasticities across sectors. We follow the estimating strategy presented in [Aguiar and Bils \(2015\)](#). Denoting a household by n , [Aguiar and Bils](#) propose to estimate expenditure elasticity of sector s , η_s , as

$$\ln \left(\frac{x_{st}^n}{\bar{x}_{st}} \right) = \alpha_{str} + \eta_s \ln E_t^n + \Gamma_s Z^n + u_{st}^n, \quad (19)$$

where x_{st}^n is the expenditure in sector s goods from household n during quarter t , \bar{x}_{st} denotes the average expenditure in s across households during quarter t , α_{str} denotes quarter-and-year-of-interview and region-of-residence fixed effects, E_t^n denotes total quarterly expenditure of household n , Z^n is a vector of demographic controls (dummies for age bins, number of earners, and household size) and u_{st}^n is an error term. [Aguiar and Bils \(2015\)](#) argue that this specification “provides a tractable framework to address for the mis-measurement in the CEX” in which “respondent’s errors (...) are scaled up by their level of expenditures.”²⁶ To address potential measurement error in sectoral expenditures that would accumulate in the measure of total quarterly expenditure, E_t^n , we follow [Aguiar and Bils](#) and instrument total expenditures with total yearly income after taxes, and quintile dummies for the household’s income group. The rationale is that “total expenditure reflects permanent income and thus will be correlated with current income.”

Column (3) in [Table 3](#) reports the estimated expenditure elasticities η_s from [Equation \(19\)](#). We find that “Education and Health Care” and “Arts, Entertainment, Recreation and Food Services” are the most expenditure-elastic sectors. Conversely, “Construction” and “Agriculture” appear to be the least expenditure-elastic. The range in the value of expenditure elasticities that we find is similar to the range that [Aguiar and Bils \(2015\)](#) find—with demand specified over final expenditures in their case. Similar to us, they find that the most expenditure-elastic sector is education and that food at home is the second least expenditure-elastic sector (after tobacco). Finally, we note that these estimates correspond to the income elasticities reported in the [Figures 1 and 2](#) of the Introduction.

Dartmouth Atlas data. Since the CEX reports whether a household member receives Medicare or Medicaid, we can impute expenditures to these household members in a similar way to K-12. For financial services, we assume an expense ratio of 90 basis points across all household pension funds and securities ([French, 2008](#)). See [Appendix B.2](#) for more details and a comparison of the estimates with and without imputation.

²⁵In practice, this imputation method generates more conservative estimates of the expenditure elasticity for the Education and Health Care sector (it reduces the estimate by slightly less than 10%), while it hardly affects the estimates for other sectors. See [Appendix B.2](#) for the estimates without imputation.

²⁶This differencing strategy cancels out log-linearly additive errors that are good- and time-specific. Note also that the term α_{str} controls for the effect of changing prices.

Table 3: Estimated Expenditure and Demand Elasticities

Sectors	Nonhomothetic CES ²		Reduced-Form ³
	σ	ε_s	η_s
	(1)	(2)	(3)
Education and Health Care		1.80 (0.07)	1.55 (0.01)
Arts, Entertainment, Recreation and Food Services		1.39 (0.03)	1.30 (0.01)
Finance, Professional, Information, other services (excl. gov't)		1.26 (0.02)	1.22 (0.01)
Government ¹		1.00 (-)	1.06 (0.01)
Manufacturing		0.69 (0.03)	0.88 (0.01)
Retail, Wholesale Trade and Transportation		0.61 (0.03)	0.82 (0.01)
Construction		0.48 (0.04)	0.75 (0.01)
Agriculture		0.30 (0.06)	0.63 (0.01)
Elasticity of Substitution	0.45 (0.05)		

Notes: Standard errors clustered at the household level shown in parentheses. Total number of households is 7,809. 1: Gov't sector is normalized to 1 in the demand estimation. 2: η_s in column (3) is the model-implied expenditure elasticity for the sample average household expenditures at the estimated parameters $\{\varepsilon_s, \eta\}$, see Equation (16). 3: η_s in column (4) corresponds to the expenditure elasticity estimated using the Aguiar and Bils (2015) specification.

Estimation of Demand Parameters Next, we describe how we estimate our demand system. We use the same household-level expenditure data, supplemented with time series for sectoral urban prices across different US regions from the BLS (as in Comin et al., 2015). Our estimation strategy also follows Comin et al. (2015) and it is based on the generalized method of moments (GMM). To write our estimating equations, we leverage on the log-linear nature of the demand system and use it to invert the demand for one sector \hat{s} to obtain an expression for the real consumption index of household n , C_t^n , in terms of observables. We then use this expression and substitute out C_t^n in the expenditure share equations for all

$s \neq \hat{s}$ to obtain

$$\ln x_{st}^n = \ln \zeta_s^n - \frac{\varepsilon_s}{\varepsilon_{\hat{s}}} \ln \zeta_{\hat{s}}^n + (1 - \sigma) \ln \left(\frac{p_{st}^n}{p_{\hat{s}t}^n} \right) + (1 - \sigma) \left(\frac{\varepsilon_s}{\varepsilon_{\hat{s}}} - 1 \right) \ln \left(\frac{E_t^n}{p_{\hat{s}t}^n} \right) + \frac{\varepsilon_s}{\varepsilon_{\hat{s}}} \ln x_{\hat{s}t}^n. \quad (20)$$

This procedure generates a system of equations for each household n and all $s \neq \hat{s}$. Equation (20) makes clear that the system is identified up to a normalization of ε_s and ζ_s for one sector. We normalize $\varepsilon_{\hat{s}} = \zeta_{\hat{s}} = 1$ for one sector without loss of generality (in our empirical exercise we normalize to one the parameters of the Government sector).

We parametrize the taste parameter as $\ln \zeta_s^n \equiv \Gamma_s Z^n + \delta_{str} + u_{st}^n$, where Z^n denotes the same household controls as in the previous estimation (dummies for age, number of earners and family size), δ_{str} are region-of-residence \times sector and quarter-and-year-of-interview \times sector fixed effects, and u_{st}^n is an error term. After incorporating this parametrization to Equation (20), we obtain

$$\ln x_{st}^n = \Gamma_s Z^n + (1 - \sigma) \ln \left(\frac{p_{st}^n}{p_{\hat{s}t}^n} \right) + (1 - \sigma) (\varepsilon_s - 1) \ln \left(\frac{E_t^n}{p_{\hat{s}t}^n} \right) + \varepsilon_s \ln x_{\hat{s}t}^n + u_{st}^n. \quad (21)$$

This resulting system of $S - 1$ equations for $s \neq \hat{s}$ defines moments in terms of observables that we use in our estimation.²⁷ We also use the reduced-form expenditure elasticities η_s estimated using the Aguiar and Bils (2015) methodology as an additional set of moments. Taking the expression for expenditure elasticity η_s from Equation (16), we have that

$$\hat{\eta}_s = \sigma + (1 - \sigma) \frac{\varepsilon_s}{\sum_{s' \in \mathcal{S}} \bar{x}_{s'}^n \varepsilon_{s'}} + v_s, \quad (22)$$

where $\hat{\eta}_s$ is the expenditure elasticity estimated using the Aguiar and Bils (2015) method, \bar{x}_{st}^n is the sample average expenditure share in sector s , and v_s denotes an error term. In sum, we use the set of moments defined by Equations (21) and (22) to estimate the demand parameters.

To deal with potential measurement error and endogeneity concerns, we use instruments for the observed measures of household expenditures and relative prices proposed in Comin et al. (2015). As in Aguiar and Bils (2015), we use household (after-taxes) income levels and income quintiles as instruments for quarterly expenditures. These instruments capture the permanent household income, and are therefore correlated with household expenditures without being affected by transitory measurement error in total expenditures. We instrument household relative prices with a Hausman-style relative-price instrument. Each of the

²⁷Since a priori any sector can be used as a reference sector in (20), we use in our estimation all $S(S - 1)$ moments we can obtain from having each sector $s \in \mathcal{S}$ as a reference.

prices used in the relative-price instrument is constructed in two steps. First, for each sub-component of a sector, we compute the average price across regions excluding the own region. Then, the sectoral price for a region is constructed using the average region expenditure shares in each sub-component as weights. These price instruments capture the common trend in U.S. prices while alleviating endogeneity concerns due to regional shocks.²⁸

The first and second columns of Table 3 report our estimates of the nonhomotheticity parameters ε_s and the elasticity of substitution σ . We find a value for the elasticity of substitution of 0.45, implying that the eight sectors are complements. This value is very similar to the estimated elasticity of substitution of 0.42 in Goos et al. (2014), obtained using European data for 10 sectors from 1993-2010. The nonhomotheticity parameters $\{\varepsilon_s\}_{s \in \mathcal{S}}$ vary substantially across sectors, the highest being for “Education and Health Care” and the lowest, “Agriculture.” The third column reports the implied expenditure elasticity for each sector for the average household in our sample. The expenditure elasticities from our model estimates and from the Aguiar-Bils reduced are very similar.²⁹

3.3 Partial-Equilibrium Quantification

Next, we combine the estimated demand parameters with the production side described in Section 2 to conduct a partial-equilibrium quantification of our model. The goal is to assess the contribution of income and price effects to changes in sectoral value-added and wage-bill shares, taking sectoral prices and income directly from the data. This simple exercise allows us to demonstrate the relevance of the income-driven mechanism in a stripped-down setting, before turning to the general-equilibrium quantification with occupational choice of Section 4.

Calibration of the Partial-Equilibrium Model Our calibration strategy consists in matching the 1980 values of expenditure per capita, sectoral prices and sectoral composition of value added. We then shock the model economy with the 2016 values for household expenditure and sectoral prices, which are the exogenous driving forces in our partial equilibrium quantification.

We set the demand elasticity parameters $\{\varepsilon_s, \sigma\}_{s \in \mathcal{S}}$ to their estimated values in Table 3. Given the 1980 sectoral value-added deflators, $\{p_{s0}\}_{s \in \mathcal{S}}$, and the level of personal consumption

²⁸Using the average price in the U.S. excluding the own region addresses the concern of regional shocks, while capturing the common component of prices across regions. Using average expenditures in the region addresses the concern of mismeasurement of household expenditure shares in that region to the extent that the mismeasurement averages out in the aggregate.

²⁹The correlation between the two expenditure elasticities is over 0.99. We find similar demand parameter estimates using CEX data without imputation, other time periods, and when we only include region and year fixed effects. See Appendix B.2.

expenditures (PCE) per capita, E_0 , we calibrate $\{\zeta_s\}_{s \in \mathcal{S}}$ to match the sectoral distribution of value added in 1980. We also match the 1980 wage bills by taking factor intensities $\{\alpha_{js0}, \beta_{s0}\}$ from the data. To simulate the 2016 economy, we increase PCE per capita by their growth over the 1980-2016 period, and allow prices to change to their 2016 value. We also allow $\{\alpha_{jst}, \beta_{st}\}$ to change to their 2016 values. Note that we keep constant of the demand parameters, $\{\zeta_s, \sigma, \varepsilon_s\}$.³⁰

Value-added growth: model vs. data We first assess the ability of the calibrated model to replicate the changes in sectoral value-added shares observed in the data, and decompose the overall change in the contribution of income and price effects. We use Equations (14) and (17) to compute the sectoral value added that results from three exercises. The first consists in increasing total expenditure per capita E_t to match the 2016 level while increasing uniformly sectoral prices by the same factor as the Personal Consumption Expenditure (PCE) deflator, so that relative sectoral prices do not change. We call this exercise a neutral increase in expenditure. The second exercise consists in changing sectoral prices to match the 2016 level of sectoral prices relative to the PCE deflator. The third exercise simulates the effect on sectoral value added of conducting simultaneously the neutral increase in expenditure and the change in relative sectoral prices.³¹

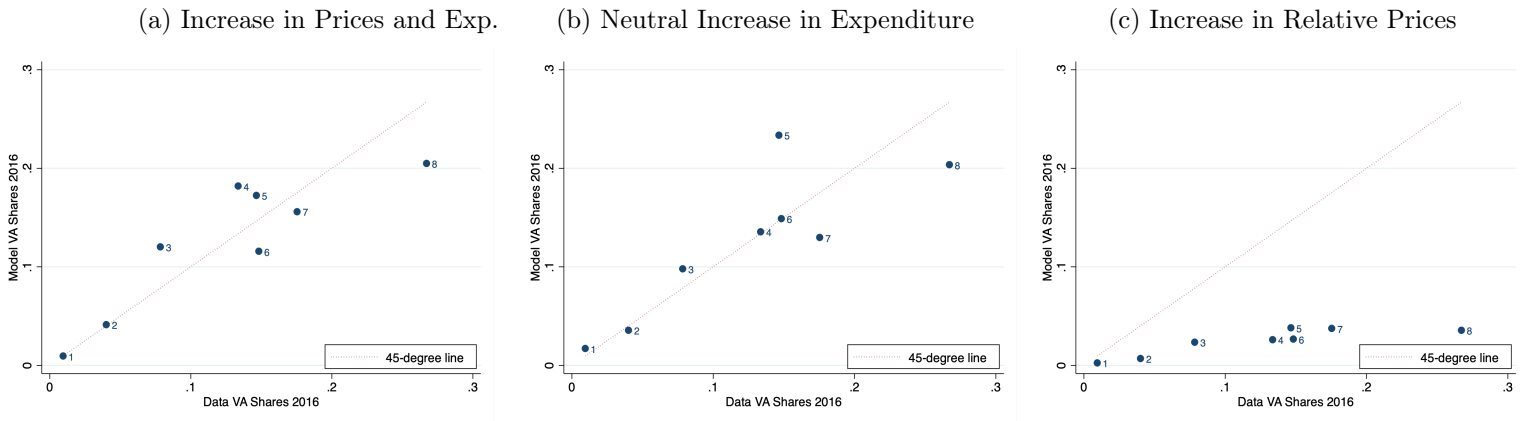
Figure 3 plots the model sectoral value-added shares for 2016 (y-axis) against the value-added shares in the data (x-axis) for the three exercises. Table 4 reports the standard deviation of the model and data value-added shares, their correlation, and their covariance normalized by the variance in the data. We begin by discussing the third exercise, which changes both total expenditure and prices. Figure 3a shows that the model accounts quite well for the sectoral composition of the economy in 2016. The dispersion of sectoral shares is similar to the data (0.069 vs. 0.082), and the correlation between sectoral shares in the data and in the model is 0.89. Consequently, the model accounts for 66% of the variation in sectoral shares we observe in 2016.

The two sectors that deviate the most from the 2016 value-added shares observed in the data are Education and Health Care (sector 4 in the plot) and Finance, Professional, Information and other services (sector 8 in the plot). In particular, the share of Education and Health Care produced by the model in 2016 is too large (18% vs. 13%), while the share of Finance, Professional, Information and other services is too small (20% vs. 26%). These discrepancies can be partly the result of biases in the measurement of sectoral prices and in

³⁰The data sources for these data are the same as in previous sections, see Appendix B for more details.

³¹For each exercise, we measure the model-implied sectoral value-added shares for 2016 as the value of sectoral value added produced by the model relative to the aggregate value added observed in the data. In doing so, we adjust the denominator to correct for the fact that from 1980 to 2016, PCE increased slightly more than nominal valued added.

Figure 3: Change in Sectoral Value-Added, 1980-2016: Data vs. Model in Partial Equilibrium



Legend: 1. Agriculture, 2. Arts, Entertainment, Recreation and Food Services, 3. Government (excl. Health and Educ.), 4. Education and Health Care, 5. Manufacturing, 6. Retail, Wholesale and Transportation, 7. Construction, 8. Finance, Professional, Information and other services.

the estimation of income elasticity parameters.³²

What are the contributions of expenditure and relative prices to the change in sectoral composition? Figure 3 and Table 4 help us answer this question. Both the neutral increase in expenditure and the change in relative prices produce model sectoral value-added shares for 2016 that are highly correlated with the data (0.81 and 0.87, respectively). The key difference between them is that the dispersion of value-added shares induced by the neutral increase in expenditure is much larger than the dispersion produced by the change in relative prices (0.074 vs. 0.013). As a result, the neutral increase in expenditure accounts for a larger fraction of the variation in sectoral shares observed in 2016 than the change in relative prices (0.67 vs. 0.12). Because we calibrated the model to match 1980 sectoral shares, we conclude that the increase in expenditure is the main driver of the observed changes in the sectoral composition of the economy.

³²Failing to adjust for quality improvements in health will result in an upward bias in measured relative price growth of the sector, and (since σ is smaller than 1) this will lead to a larger sectoral share. Reasonable adjustments to sectoral deflators can correct the discrepancy between model and data. For example, in Section 4, we show that if the price deflator of Education and Health Care grew 23% more than the PCE price deflator during the period 1980-2016, the model matches the share of Education and Health Care in the data. Similarly, despite our efforts to impute consumption of financial services for which households (especially those with higher income) do not pay directly, there are numerous financial fees, commissions and services that are not reflected in the CEX, as well as business services that may be provided by employers as part of compensation. Since those would increase the expenditures in the Finance, Professional and Information services sector of richer households, they would imply a value of the income elasticity parameter larger than the one we have estimated. Conversely, there are many services provided by the government to households that are not reflected in the CEX (e.g., defense, justice, etc.). To the extent that those are enjoyed by households regardless of their income, our estimates for the income elasticity of government services would be biased upwards. This could explain why our model predicts a larger increase in the share of government services than in the data.

Table 4: Sectoral Value Added Shares in 2016 in the Partial-Equilibrium Model

	Data	Model Changes:		
		$\{E, p_s\}$	E	$\{p_s\}$
Std. Dev.	0.082	0.069	0.075	0.013
Corr(Data, Model)		0.89	0.84	0.87
Cov(Data, Model)/Var(Data)		0.66	0.67	0.12

Decomposing Changes in wage-bill shares Substituting Equation (17) in (5), we obtain an expression for the wage bill of workers in occupation j as a function of total expenditure and prices,

$$w_{jt}X_{jt} = \sum_{s \in \mathcal{S}} \hat{\alpha}_{jst} V A_{st} = \sum_{s \in \mathcal{S}} \hat{\alpha}_{jst} \zeta_s E_t^{\sigma + \varepsilon_s} p_{st}^{1 - \sigma} P_t^{-\varepsilon_s}. \quad (23)$$

Expression (23) allows us to explore the role of nonhomotheticities and prices in the evolution of the wage-bill share across occupations. By construction, our calibration matches the shift term in the data—given that we match initial value-added shares and take factor intensities from the data. We thus focus on the share term.

The second row in Table 5 reports the share term produced by the model when we change both expenditures and relative prices from their 1980 levels to their 2016 levels. A comparison of the model-produced share term with the share contribution in the data suggests that the model accounts well for the share term in the data. To study the role of the income-driven mechanism, we generate the share term by increasing expenditures to their 2016 level while keeping relative prices at their 1980 level in the model. The results are reported in the third row. The fourth row shows that the neutral increase in expenditure accounts for 57% of the increase in the wage-bill share of low-skill occupations. For middle- and high-skill occupations, we find that it accounts for 12% and 27% of the actual changes in the wage-bill shares of these occupations.

The results presented in this section suggest that the income-driven channel plays a substantial role in generating both sectoral reallocation of value-added and labor-market polarization. Importantly, the demand parameters governing the income effects and the elasticity of substitution across sectors come from cross-sectional household estimates. Thus, they are independent from the aggregate time-series phenomena that we are trying to explain. Despite this, the partial equilibrium model does a good job in generating the polarization of labor markets, with the income-driven channel playing a significant role.

Table 5: Change in Wage Bill Shares, Partial-Equilibrium Quantification

	Occupations		
	Low	Middle	High
Total Change (Data)	0.021	-0.21	0.19
Model-generated Share Term:			
All channels (E and prices)	0.026	-0.08	0.09
Only E	0.012	-0.02	0.05
Only E /Data Total	57%	12%	27%

4 General-Equilibrium Model with Occupational Choice

So far, our analysis of the income-driven mechanism has focused on occupational wage bills in a representative-agent, partial-equilibrium setting. To provide a richer account of labor-market polarization, we introduce heterogeneous households facing an occupational choice à la Roy. This extension of our model generates separate predictions for hours worked and relative wages across occupations. Additionally, by making sectoral prices and household expenditures endogenous, the general equilibrium nature of the model can shed light on the economic primitives through which the income-driven channel generates polarization.

4.1 Model

Our model uses the production structure and preferences presented in Sections 2.1 and 3.1. Since this section introduces heterogeneous households that self-select into different occupations, we briefly show how we extend our framework to incorporate this heterogeneity. We then present the occupational choice problem and define the competitive equilibrium.

4.1.1 Production

The setting is identical to Section 2.1, except that labor inputs in the production function are now expressed in terms of efficiency units, rather than total hours. This change allows us to account for household heterogeneity in labor supply. That is, the representative firm in each sector $s \in \mathcal{S}$ operates with the technology defined in Equation (1)

$$Y_{st} = A_{st} \left(\prod_{j \in \{L, M, H\}} \tilde{X}_{jst}^{\alpha_{jst}} \right)^{\beta_{st}} K_{st}^{1-\beta_{st}}, \quad (24)$$

where the only difference is that \tilde{X}_{jst} denotes the number of efficiency units of labor employed in occupation j sector s , and year t . All derivations can be done in an analogous way to Section 2 and are relegated to Appendix C.

4.1.2 Household preferences, endowments and demographics

There is a continuum of mass 1 of households indexed by $h \in \mathcal{H} \equiv (0, 1)$. Each household is endowed with one unit of labor and K_{ht} units of capital. Households face an occupational choice (which we discuss in more detail below). They choose which occupation (high-, middle- or low-skill) to inelastically supply their unit of labor. Household income is composed of the labor income plus the rental income accrued from the owned capital. In our baseline exercise, we assume all capital is evenly distributed across households.³³

We maintain the static nature of our model from Section 3, and assume that household expenditure E_{ht} equals household income period by period. Each household maximizes utility U_{ht} , defined by the same nonhomothetic CES aggregator as in Equation (12),

$$\sum_{s \in \mathcal{S}} (U_{ht}^{\varepsilon_s} \zeta_s)^{\frac{1}{\sigma}} c_{hst}^{\frac{\sigma-1}{\sigma}} = 1, \quad (25)$$

subject to the household budget constraint $E_{ht} \geq \sum_{s \in \mathcal{S}} p_{st} c_{hst}$. Note that all demand parameters, $\{\zeta_s, \varepsilon_s, \sigma\}_{s \in \mathcal{S}}$ are the same across households and constant over time. This implies that household heterogeneity in consumption choices is solely driven by differences in household income at any given point in time. Since household demand can be derived as in Section 3.1, we also relegate these derivations to Appendix C.

Occupational Choice A key difference with the model presented in Section 3.1 is that now households face an occupational choice which yields an endogenous labor supply at the occupational level. Household h draws a vector $(\eta_{hH}, \eta_{hM}, \eta_{hL})$ of potential efficiency units to be supplied in each occupation. Given the vector of wages per efficiency unit in the three occupations $(\tilde{w}_{Ht}, \tilde{w}_{Mt}, \tilde{w}_{Lt})$, household h chooses the occupation that maximizes labor income,

$$\operatorname{argmax}_{j \in \{L, M, H\}} \{\eta_{hj} \tilde{w}_{jt}\}. \quad (26)$$

To quantitatively evaluate the model, we assume that efficiency units for each occupation are independently drawn from lognormal distributions. We denote the mean and standard deviation of $\log \eta_j$ by μ_j and χ_j .³⁴ We interpret these idiosyncratic draws of efficiency units

³³In our robustness section, we explore an alternative setting where capital is owned by households working in high-skill occupations.

³⁴Lognormal distributions have been used in this context, see for example Bárány and Siegel (2018) and Cerina et al. (2017). We note that assuming a Fréchet distribution (or a multi-variate Fréchet in the max-

as a reduced-form mapping from ability to worker productivity across occupations. Since this formulation assumes that all households are ex-ante identical, it abstracts from an array of pre-existing differences across groups in the population, e.g., gender and race (Hsieh et al., 2013) that are critical to account for labor supply decisions at the household level. We abstract from these ex-ante differences because our goal is to parsimoniously generate a labor supply at the occupation level.

4.1.3 Competitive Equilibrium

A competitive equilibrium is defined by a sequence of prices $\{\{p_{st}\}_{s \in \mathcal{S}}, \tilde{w}_{Lt}, \tilde{w}_{Mt}, \tilde{w}_{Ht}, r_t\}_{t=0}^{t=T}$, allocations $\{\{c_{hst}\}_{s \in \mathcal{S}, h \in \mathcal{H}}\}_{t=0}^T$, capital holdings $\{\{K_{ht}\}_{h \in \mathcal{H}}\}_{t=0}^T$ and household occupational choices for each household h such that:

1. Each household maximizes utility (25) subject to the budget constraint $\sum_{s \in \mathcal{S}} p_{st} c_{hst} = E_{ht}$ where $E_{ht} = \max_{j \in \{L, M, H\}} \{\eta_{hj} \tilde{w}_{jt}\} + r_t K_{ht}$.
2. Firms maximize profits taking prices as given, $\max_{\{\tilde{X}_{jt}, K_{st}\}} p_{st} Y_{st} - \sum_j \tilde{w}_{jt} \tilde{X}_{jt} - (r_t + \delta) K_{st}$, where it is understood that $j \in \{L, M, H\}$.
3. All markets clear. In particular, aggregate labor supply equals aggregate demand for each occupation,³⁵ aggregate capital demand equals aggregate capital supply, and goods markets clear, $\int_0^1 c_{hst} dh = Y_{st}$, for all $s \in \mathcal{S}$.

4.2 Model Calibration

To quantitatively evaluate our model, we need to specify the value of the preference parameters, $\{\zeta_s, \varepsilon_s, \sigma\}_{s \in \mathcal{S}}$, sectoral technology parameters $\{\alpha_{st}, \beta_{st}, A_{st}\}_{s \in \mathcal{S}, t \in \{1980, 2016\}}$, the aggregate supply of capital K_t in $t \in \{1980, 2016\}$, the depreciation rate δ , and the parameters of the lognormal distributions over which labor efficiency units $\{\eta_j, \chi_j\}_{j \in \{L, M, H\}}$ are drawn from. Our calibration strategy consists of two parts. First, we set the values of all model parameters in 1980 to exactly match the 1980 economy. Second, we change some of model parameters to match key moments in 2016, while keeping the rest of the parameters fixed at their initially calibrated values. We outline our targets and procedure for the calibration in the main text.

stable family as described in Lind and Ramondo, 2018) in this setting would have the counterfactual prediction that average wage per worker is equalized across occupations. Authors that have used the Fréchet distribution in similar settings need to resort to unobserved costs or worker attributes, see e.g., Galle et al. (2017).

³⁵Aggregate demand follows from Equation (6). See Equation (21) in Appendix D for aggregate labor supply.

Appendix D presents a more detailed explanation and discussion of the calibration procedure. The values that result from the calibration procedure for the aggregate parameters are reported in Table 6.

Calibration of the 1980 Economy We set the preference parameters σ , and $\{\varepsilon_s\}_{s \in \mathcal{S}}$ to their estimated values from Table 3 and the depreciation rate δ to 10%. We then calibrate the rest of the model parameters, $\{\alpha_{s1980}, \beta_{s1980}, A_{s1980}, \zeta_s\}_{s \in \mathcal{S}}$, the aggregate capital stock K_{1980} and the distribution over draws of efficiency units to match the 1980 values of the following moments:³⁶

- The aggregate nominal value added per capita, $\sum_{s \in \mathcal{S}} VA_{s1980}$,
- the sectoral distribution of value added, $\{VA_{s1980}\}_{s \in \mathcal{S}}$,
- sectoral prices, $\{p_{s1980}\}_{s \in \mathcal{S}}$,
- the interest rate r_{1980} ,
- the wages per worker in high- and low-skill occupations relative to middle-skill occupations,
- the wage bills in each sector accrued by workers in each of the three occupation categories, $\{\tilde{w}_{js1980} \tilde{X}_{js1980}\}_{s \in \mathcal{S}, j \in \{L, M, H\}}$.

We next discuss our calibration procedure. Firm optimization implies that the factor shares by occupation and sector directly correspond to the exponents in the Cobb-Douglas production function, $\{\alpha_{sj1980}, \beta_{s1980}\}_{s \in \mathcal{S}, j \in \{L, M, H\}}$ (as shown in Equation 2).³⁷ Given sectoral value added and the sectoral capital shares, the rental cost of capital (interest rate plus depreciation) pins down the level of capital per capita used in each sector (Equation 3 in Appendix C). Adding across sectors, this corresponds to the aggregate capital stock per capita, K_{1980} (Equation 7).

We calibrate the parameters governing the lognormal distributions generating the efficiency draws as follows. First, since the definition of an efficiency unit for each occupation is arbitrary, we can normalize the average level of productivity for each occupation $\{\mu_j\}_{j \in \{L, M, H\}}$. We set μ_j equal to a common value μ_{1980} for all occupations to match the level of aggregate nominal value added per capita. Since labor earnings are homogeneous of degree one in efficiency draws $\{\eta_j\}_{j \in \{L, M, H\}}$, only relative efficiency draws matter for occupational choice.³⁸ This allows us to normalize the variance of draws in one occupation

³⁶See Appendices B and C for more details on our data sources and parameter calibration.

³⁷See Equations (2) and (4) in Appendix C for the expressions with efficiency units in the production function.

³⁸This is a well-known result in this literature. It is used, among others, by Autor and Dorn (2013).

(we choose low-skill). We then calibrate the variance of the productivity distributions across middle- and high-skill occupations, $\{\chi_j\}_{j=\{M,H\}}$, to match: (1) the wage-bill shares across occupations, and (2) average wages per worker of high- and middle- relative to low-skill occupations. Note that by matching the wage bills and the relative wages in 1980, we also match the employment shares across occupations in 1980.

Finally, we use the dependence of producers' cost function on sectoral TFP (see Equation 5 in Appendix C) and set sectoral TFP levels in 1980, $\{A_{s1980}\}_{s \in \mathcal{S}}$, to match the observed sectoral price deflators. Given the distribution of household income and the preference parameters, $\{\varepsilon_s, \sigma\}_{s \in \mathcal{S}}$, we set the taste parameters $\{\zeta_s\}_{s \in \mathcal{S}}$ so that the model-generated aggregate sectoral consumption matches the observed sectoral value added in 1980.

Calibration of the 2016 Economy The second part of our calibration consists in setting the parameters that change from 1980 to 2016 to their new values. We assume that the preference and taste parameters $\{\zeta_s, \varepsilon_s, \sigma\}_{s \in \mathcal{S}}$, the depreciation rate δ , and the dispersion of productivities in each occupation, $\{\chi_j\}_{j=\{L,M,H\}}$, do not change over the horizon we study. We allow changes in the common mean of the lognormal distribution for the three occupations, μ_{2016} , the sectoral TFP levels $\{A_{s2016}\}_{s \in \mathcal{S}}$, the factor intensities in each sector $\{\alpha_{sj2016}, \beta_{s2016}\}_{s \in \mathcal{S}, j=\{L,M,H\}}$, and the level of capital, K_{2016} . The new values for these parameters are set to match the following moments:

- The 2016 share of sectoral nominal value added, $\{VA_{s2016}/VA_{2016}\}_{s \in \mathcal{S}}$,
- the wage-bill shares by sector and occupation in 2016, $\{\tilde{w}_{js2016}\tilde{X}_{js2016}\}_{j \in \{L,M,H\}}$,
- the interest rate r_{2016} ,
- the increase in aggregate nominal expenditures per capita from 1980 to 2016, E_{2016}/E_{1980} ,
- the increase in the Fisher price index for personal consumption expenditures from 1980 to 2016.

We match these targets as follows. The factor intensities $\{\alpha_{sj2016}, \beta_{s2016}\}_{s \in \mathcal{S}, j=\{L,M,H\}}$ are set to match the 2016 labor and occupation shares in sectoral value added (as implied by the firms' optimality conditions). Given r_{2016} , we obtain the aggregate level of capital per capita, K_{2016} , aggregating the capital demand across sectors, which we compute combining sectoral labor shares and value added (as in the 1980 calibration). To match sectoral value added, we proceed as follows. Let $\hat{A}_{s2016} \equiv A_{s2016}/A_{2016}$ be the TFP level of sector s relative to the aggregate TFP in the economy, A_{2016} . We calibrate simultaneously $\{\hat{A}_{s2016}\}_{s \in \mathcal{S}}$, A_{2016} , and μ_{2016} to match the 2016 sectoral shares of nominal value added, the growth in aggregate expenditures per capita, and the growth in the Fisher price index of personal consumption

Table 6: Calibrated Aggregate Parameters

	1980	2016
	Time-Invariant	
Variance of Efficiency Units		
χ_L	1.00	1.00
χ_M	0.81	0.81
χ_H	1.19	1.19
Depreciation, δ	0.10	0.10
	Time-Varying	
Capital pc, K_t	0.35	2.91
TFP, A_t	1.00	1.18
Labor Prod., μ_t	0.01	0.04

expenditures from 1980 to 2016. Note that by matching sectoral shares and the growth in total expenditure per capita, we also match the sectoral distribution of value added.

The assumption that preference parameters $\{\zeta_s, \varepsilon_s, \sigma\}_{s \in \mathcal{S}}$ and the variances of labor productivity across occupations $\{\chi_j\}_{j=\{L,M,H\}}$ remain constant to their 1980 values introduces two significant differences between the calibrations for 1980 and 2016. First, we do not target the sectoral prices in 2016. Second, the calibration procedure is designed to match the occupational wage-bill shares in 2016, but it does not require that the model matches the distribution of hours worked or the relative wages across occupations in 2016.

4.3 Model Quantification Results

Untargeted moments Before analyzing the drivers of polarization, we study how well the model fits untargeted moments such as the distribution of hours worked and relative wages in 2016, and sectoral price growth from 1980 to 2016. Tables 7 and 8 report the relevant data moments and model outcomes. The first three rows in Table 7 report the model-implied changes in wage-bill and employment shares vis-à-vis the data. By construction, our calibration matches the change in wage-bill shares and the contributions of the shift and share terms. Despite not being a calibration target, the model captures well the polarization in hours worked across occupations. The third line in the right panel shows that the change in hours worked across occupations produced by the model are 82%, 91% and 93% of those observed in the data for low-, middle- and high-skill occupations. A shift-share decomposition (as in Equation 11 in Section 2) reveals that the slight discrepancy between model and data is due to the shift term and that the model’s implied share term is remarkably close to that in the data (0.039 vs. 0.037 for high-skill, -0.064 vs. -0.061 for middle-skill, and 0.025 vs.

Table 7: Baseline Model Quantification

	Change Wage Bill Sh.			Change Employ. Sh.		
	Low	Middle	High	Low	Middle	High
Data	0.021	-0.209	0.188	0.034	-0.165	0.131
Model	0.021	-0.209	0.188	0.028	-0.15	0.122
Model/Data(%)	100%	100%	100%	82%	91%	94%
Decomposition of Model:						
Shift	0.006	-0.151	0.13	0.004	-0.089	0.085
Share	0.015	-0.058	0.058	0.024	-0.061	0.037
Share/Model(%)	73%	28%	31%	86%	41%	31%
Contribution of Different Channels:						
Shift, E Shock	0	0	0	0	0	0
Shift, Biased-Tech. Shock	0.006	-0.151	0.13	0.004	-0.089	0.085
Share, E Shock	0.019	-0.005	0.066	0.018	-0.052	0.034
Share, Biased-Tech. Shock	-0.008	-0.055	-0.054	0.005	0.003	-0.008
Share, E Shock/Model (%)	90%	2%	35%	64%	34%	28%

Notes: Numbers in blue are targeted in the calibration. The calibration strategy also matches 1980 wage bill and employment shares. See the main text for further discussion.

0.024 for low-skill occupations). Table 8 shows that our model also generates polarization in relative wages that is in line with the data. If anything, the model somewhat overpredicts the relative wage in 2016 of low- to middle-skill occupations (0.86 vs. 0.8 in the data), and high- to middle-skill occupations (1.57 vs. 1.53 in the data).

The sectoral prices that our calibrated model generates to fit the 2016 sectoral value-added shares are not very different to those in the data. Figure A.1a in Appendix A contains a scatter plot of the growth in sectoral prices from 1980 to 2016 in the model and the data. Their correlation is 0.63. The standard deviation of growth rates is 1.44 in the data, and 1.65 in the model. The main deviations between data and model are in three sectors: “Education and Health Care,” “Government,” and “Finance, Professional, Information and Other services.” In the first two, model prices grow by less than in the data, while in the latter they grow by more. One possible interpretation of these differences in sectoral price growth between model and data in “Education and Health Care” is that it partly reflects unmeasured quality improvements (Jaimovich et al., 2019). In the “Finance, Professional and Information services” and “Government” sectors, instead, the discrepancy in prices may reflect the fact that estimates of the income elasticity parameters based on the CEX may be somewhat inaccurate because they missed part of the actual household consumption, as further discussed already in footnote 32.

Table 8: Baseline Model Quantification for Relative Wages

	Year	w_L/w_M	w_H/w_M
Data	1980	0.74	1.24
	2016	0.8	1.49
Model	1980	0.74	1.24
	2016	0.86	1.57
Contribution of Different Channels:			
Only E	2016	0.79	1.32
Only Biased Tech.	2016	0.8	1.46
Only E/Model (%)		46%	29%

Notes: Numbers in blue are targeted in the calibration. The last line reports the percent of the total increase in relative wages generated by the E shock plus half of the covariance when both E and biased-technology shocks are included.

In sum, we conclude that our model does a remarkable job in generating the polarization of key untargeted variables such as hours worked and wages, and produces a reasonable evolution for relative sectoral prices.

4.3.1 Untangling Polarization Mechanisms

We next use our model to shed light on the drivers of the shift and share components of polarization in wage bills, hours worked, and relative wages across occupations.

Grouping parameter changes In the spirit of our partial equilibrium quantification from Section 3.3, we group the simulated changes in parameters from 1980 through 2016 in two distinct “shocks.” The first shock consists in the increases in aggregate TFP, A_{2016} , and in the common mean of the labor productivity distribution, μ_{2016} . This shock resembles the neutral increase in expenditure from the partial equilibrium quantification in Section 3.3 because it is the main driver of growth in nominal expenditures and the PCE deflator. Consequently, we call it the expenditure shock, or E shock for short. The second shock consists of the changes in relative sectoral TFP, $\{\hat{A}_{s2016}\}_{s \in \mathcal{S}}$; factor intensities, $\{\alpha_{js2016}, \beta_{s2016}\}_{s \in \mathcal{S}}$; and the level of capital, K_{2016} .³⁹ This shock is the natural counterpart to the exogenous changes in relative sectoral prices from the partial equilibrium exercise, and because it also includes exogenous changes in factor intensity, we label it the biased-technology shock.⁴⁰ It is

³⁹We group the change in the capital stock into the second shock because, for given nominal value added and factor intensities (which we target in our calibration), the capital stock only impacts the rental rate and hence sectoral prices.

⁴⁰The E shock induces growth in nominal expenditures by 3.53 and in the PCE deflator by 1.18, while the biased-technology shock induces growth in nominal expenditures by 0.16 and in the PCE deflator by 0.33. The variance of sectoral price growth induced by the biased-technology shock is 0.94, while for the E shock

important to note that because prices and wages are endogenous in general-equilibrium, the biased-technology shock has an effect (albeit small) on aggregate expenditures and on the PCE deflator, while the E shock affects sectoral prices. As we discuss below, this latter effect is not negligible, and it contributes to the overall importance of the E shock for polarization.

Drivers of polarization What are the contributions of the E and biased-technology shocks to labor-market polarization? The bottom panels of Tables 7 and 8 answer this question. The first observation to note is that, by construction, the E shock has no effect on shift components. The contribution of the E shock to polarization, therefore, operates only through the changes it induces in the sectoral composition of value added/employment.⁴¹ Table 7 shows that the E shock has caused 90% and 35% of the increase in the wage-bill shares of low- and high-skill occupations, while it is responsible for only 2% of the decline in the wage-bill share of middle-skill occupations. The E shock has driven slightly smaller fractions of the increase in the shares of hours worked by low- and high-skill occupations (64% and 28%), while it is responsible for 35% of the decline in the share of hours worked in middle-skill occupations. The E shock is also key to understanding the changes in the observed relative wages across occupations from 1980 through 2016. In particular, Table 8 shows that it is responsible for 46% of the increase in the wage of low- relative to middle-skill occupations, and for 29% of the increase in the wage of high- relative to middle-skill occupations.

To understand the mechanism by which the E shock generates sectoral reallocation and polarization, it is useful to compare the results in the partial- and general-equilibrium quantification. A comparison of Tables 5 and 7 reveals that the contribution to polarization of the income-driven channel is greater in general equilibrium than in partial equilibrium. For example, in partial equilibrium, the income-driven channel causes an increase in the wage-bill share of high-skill occupations of 5 percentage points, while in general equilibrium it causes an increase of 6.6 percentage points. These differences arise because sectoral prices are endogenous in general equilibrium. In both the partial- and general-equilibrium quantifications, the increase in aggregate expenditures directly affects polarization through the effect of nonhomotheticities in the sectoral reallocation of value added. Since the E and neutral expenditure shocks are calibrated to match the actual increase in aggregate expenditures in both exercises, and the initial expenditure elasticities of demand and initial factor intensities are also the same, the direct effect of the income-driven channel is also the same in the partial- and general-equilibrium quantifications.⁴²

it is 0.65.

⁴¹The reason is that shift terms are proportional to sectoral changes in factor intensities, which are only affected by the biased-technology shock (e.g., see Equation 8).

⁴²Having a distribution of consumers with heterogeneous expenditures could, a priori, introduce a difference

In general equilibrium, however, the income-driven channel has an indirect effect on reallocation and polarization through its impact on sectoral prices. As the income-driven channel polarizes the distribution of wages across occupations, the relative prices of goods and services from sectors that are more intensive in high- and low-skill occupations increase. Since sectors are complements, this causes additional reallocation of value added towards these sectors and amplifies the direct effect of the income-driven channel on polarization. These effects can be appreciated in Figure A.1b in appendix A, where we observe that the E shock produces faster price growth in the expenditure-elastic sectors. The correlation between these two variables is 0.8. In contrast, the correlation between the expenditure elasticity of demand and the growth in prices induced by the biased-technology shock is -0.12 (see Figure A.1c). As a result, the contribution of the technology-biased shock to the share term for high- and low-skill occupations is negative.

Of course, the biased-technology shock has an important effect on polarization through the shift component. Consistent with the literature on biased technological change, routinization, and capital-skill complementarities, we find that the biased-technology shock accounts for the bulk of the decline in the wage-bill share of middle-skill occupations, and for 65% of the increase of high-skill occupations. Taken together, these results show that the income-driven channel and biased-technology forces have simultaneously carved both ends of the US labor market, making it polarize over the past 40 years. While the income-driven channel played a dominant role in shaping the outcomes for low-skill occupations, the biased-technology forces drove the effect on middle-skill occupations, and both contributed to the outcomes for high-skill occupations, with a 35-65 split in favor of the biased-technology forces.

4.4 Polarization by Subperiods and Robustness Checks

After having established the quantitative significance of the income-driven mechanism, we analyze our mechanism by subperiods, and also show that our quantitative findings are robust to extending the model to include a CES aggregator across occupations, trade or assigning capital to high-skilled only.

Analysis by Subperiods We explore how the importance of the drivers of polarization may have changed over the last four decades by splitting the interval we have studied in two subperiods: 1980-2000 and 2000-2016. We redo our analysis separately for each subperiod, so that the model matches the evolution of the wage-bill shares, and the wage-bill shift-share

between the direct effect of the income-driven channel to polarization in the general- and partial-equilibrium quantifications. We have found, however, that heterogeneity of the income distribution across households does not affect the overall magnitude of the income-driven channel.

decomposition in both subperiods.⁴³ Tables A.2 and A.6 in the appendix show that the model captures well the evolution of hours worked across occupations and relative wages per worker. The contribution of the share term in general, and of the E shock in particular, is significant across all polarization outcomes in both subperiods. We find, however, that the share term and the E shock are quantitatively more important in the first subperiod. For example, the E shock increases the share of low- and high-skill occupations in the total wage bill by 0.0126 and 0.044 during the 1980-2000 period, but only by 0.005 and 0.023 during the 2000-2016 period. This larger effect of the E shock during the first subperiod is due to faster growth in personal consumption expenditures in this period (210%) relative to the second (65%). As a result, we calibrate a larger E shock in the first period that induces greater changes in value-added shares towards high- and low-skill intensive sectors away from middle-skill intensive sectors.

CES Production Function Our baseline exercise assumes that the elasticity of substitution across different occupations is one. We relax this assumption and consider the following sectoral production functions

$$Y_{st} = A_{st} K_{st}^{1-\beta_{st}} \left(\sum_{j \in \{L, M, H\}} \alpha_{jst}^{\frac{1}{\nu}} \tilde{X}_{jst}^{\frac{\nu-1}{\nu}} \right)^{\beta_{st} \frac{\nu}{\nu-1}}, \quad (27)$$

where ν is the elasticity of substitution between occupations.⁴⁴ In our simulations, we set $\nu = 1.42$, the baseline value used in Buera et al. (2018) and the point estimate obtained by Katz and Murphy (1992).⁴⁵ We calibrate $\{\alpha_{jst}\}$ to match the wage-bill share of occupation j in sector s at time t . Table A.4 reports our results. Note that the contribution of the share term and E shock to the changes in wage-bill shares is the same as in our baseline model.⁴⁶ We find that the predicted employment shares and the contribution of the E shock both for wage-bill and employment shares are very similar to our baseline estimates. For example, it accounts for 91%, 2%, and 35% of the change in wage-bill shares for low-, middle- and high-skill occupations. Likewise, we find a similar effect for relative wages (see Table A.6).

⁴³We keep constant the preference parameters throughout this exercise. In the second subperiod, we recalibrate the dispersion in productivities for middle- and high-skill occupations, χ_M and χ_H , so that we can exactly match the relative wages per worker in the subperiod that starts in 2000.

⁴⁴We impose the normalization that $\sum_{j \in \{L, M, H\}} \alpha_{jst} = 1$.

⁴⁵This estimate assumes an aggregate production function and is based on only two skill levels. We have also conducted our own estimation of the elasticity based on more sectors and three skill levels, using variation across local labor markets in the US. When instrumenting labor supply with a Bartik-style instrument for immigrants (with initial shares of immigrants by location and aggregate migration trends), we find an elasticity of substitution of 1.2, but we cannot reject the null of 1 (i.e., Cobb-Douglas).

⁴⁶This is the case because, as anticipated in footnote 13, our baseline specification captures flexibly time variation in factor intensities which (together with the sectoral composition of value added) fully determine the share term.

Accounting for International Trade Our baseline analysis assumed a closed economy. This section extends our model to account for sectoral trade. We account for this possibility in a parsimonious way through the use wedges. Rearranging the identity that aggregate production equals domestic consumption plus net exports, $p_{st}Y_{st} = p_{st}C_{st} + NX_{st}$, we have that

$$p_{st}C_{st} = (1 - \tau_{st})p_{st}Y_{st}, \quad (28)$$

where $\tau_{st} \equiv \frac{NX_{st}}{p_{st}Y_{st}}$ captures the wedge between domestic production and consumption. If net exports are positive, trade magnifies domestic demand (i.e., enhances the changes in production and labor demand considered in the baseline model), while if net exports are negative, domestic aggregate demand is dampened. We compute the wedges $\{\tau_{st}\}$ as the value-added content of US net exports over US value-added production (see Appendix E for the details).

Tables A.3 and A.6 present the results of simulating the E and biased-technology shocks where the later now includes the changes in trade wedges. The main take away is that the role played by the E shock in the polarization of labor markets remains when taking into account the effect of international trade. For example, it accounts for 73%, 28% and 31% of the change in the wage-bill shares for low-, middle- and high-skill occupations. The intuition for this result is that the change in net exports is only significant for manufacturing but, even there, its magnitude relative to overall value added is small.

Allocation of capital holdings We study the robustness of our findings to an alternative assumption on the distribution of capital ownership across occupations. Given that most capital is owned by the richest households, a more realistic assumption about the distribution of capital ownership than the one made in the baseline exercise would be that capital is only owned by the workers in high-skill occupations.⁴⁷ Tables A.5 and A.6 show that our quantitative assessment about the drivers of polarization are robust to the rule used to allocate capital across occupations (e.g., the E shock accounts for 74%, 27% and 31% of the change in wage-bill shares for low-, middle- and high-skill occupations).

5 Extensions

Income growth and nonhomotheticities in demand are ubiquitous. Therefore, the income-driven mechanism emphasized in our model should be relevant in other periods of history and other countries. In this section, we explore this possibility. First, we study the role of sectoral reallocation in generating polarization in other advanced economies during the

⁴⁷We assume that workers do not take into account the assignment of capital to high-skill occupations when making their occupational choice.

Figure 4: US vs other countries' change in value-added shares 1980-2016, $\Delta(VA_{st}^c/VA_t^c)$



Notes: A dot represents the 2016-1980 change in value-added shares in one sector for a country (Austria, Spain, Finland, France, Germany, Italy, Japan, Netherlands, UK and US). See Appendix B.3 for details.

1980-2016 period. Second, we move the time horizon of our analysis for the US backwards and forward. Specifically, we study the drivers of labor market polarization during the period 1950-1980, and conclude by forecasting the evolution of labor-market outcomes over the next fifteen years.

5.1 Polarization in Other Advanced Economies

Labor-market polarization is a pervasive phenomenon across advanced economies. Europe and Japan have experienced similar processes to the US.⁴⁸ In this subsection, we explore the role of sectoral reallocation and the income-driven channel in generating labor-market polarization in other advanced economies. One challenge we encounter when studying other economies is that data limitations are greater than in the US.⁴⁹ For this reason, we explore the role of the income-driven channel by showing the relevance of the share term for the evolution of wage-bill and employment shares, and conclude with a quantification of the income-driven channel using the partial-equilibrium framework developed in Section 3.

We begin by showing that changes in sectoral value-added shares are highly correlated across advanced economies. Figure 4 plots the change in value-added shares in the US from 1980 to 2016 (x-axis) and the change in value-added shares in other advanced countries over

⁴⁸For example, Goos et al., 2014 document polarization in hours worked for various European countries during the period 1990-2010. Ikenaga and Kambayashi (2016) provide evidence of labor-market polarization for Japan.

⁴⁹See Appendix B.3 for details about the data used in this section.

the same period (y-axis). Our cross-country data comes from EUKLEMS and it includes the following countries: Austria, Spain, Finland, France, Germany, Italy, Japan, the Netherlands, and the UK. Figure 4 shows that the patterns of sectoral reallocation are very similar across advanced economies. In particular, the (pooled) correlation of the changes in sectoral shares with the US is 0.88.

Next, we explore shift-share decompositions of wage-bill and hours-worked shares. We begin by discussing the wage-bill shift-share decomposition. Using the EU Labour Force Survey (LFS) and the Statistics on Income and Living Conditions (EU-SILC) micro datasets supplemented with EUKLEMS, we compute factor shares in 2016 by sector, occupation, and country, $\hat{\alpha}_{jstc}$.⁵⁰ Combining factor shares for 2016, with the data on value-added shares changes between 1980 and 2016 (shown in Figure 4), we can obtain the share term of a shift-share decomposition for the wage bills, as the following equation shows

$$\frac{WB_{jtc}}{WB_{tc}} - \frac{WB_{j0c}}{WB_{0c}} = \underbrace{\sum_{s \in \mathcal{S}} \left(\frac{\hat{\alpha}_{jstc}}{\beta_{tc}} - \frac{\hat{\alpha}_{js0c}}{\beta_{0c}} \right) \frac{VA_{s0c}}{VA_{0c}}}_{Shift} + \underbrace{\sum_{s \in \mathcal{S}} \frac{\hat{\alpha}_{jstc}}{\beta_{tc}} \left(\frac{VA_{stc}}{VA_{tc}} - \frac{VA_{s0c}}{VA_{0c}} \right)}_{Share}. \quad (29)$$

We find that the share term in the sample of European economies is similar to the US. For low-skill occupations, the median share value in Europe is 0.014 vs. 0.015 in the US. For middle- and high-skill occupations, they are -0.045 and 0.041 vs. -0.058 and 0.058 in the US.⁵¹ This exercise shows that the share component has played a similar role to the US in the polarization of European labor markets.

Since our microdata does not go back to 1980, we cannot directly measure the initial occupation factor intensities, α_{js0c} , which we need to compute the shift term in Equation (29). We overcome this limitation by assuming that the growth rate of occupation intensities between 1980 and 2016 is the same as in the US. That is, denoting the growth rate of US occupation intensities by $g_{jsUS} = (\alpha_{js2016US} - \alpha_{js1980US})/\alpha_{js1980US}$, we assume that $\alpha_{js1980c} = \alpha_{js2016c}/(1 + g_{jsUS})$. This assumption seems sensible since a key driver of factor intensities is technology, and technology has changed similarly in Europe and the US over this time period. With this assumption, we can compute the remaining terms in Equation (29) since we observe sectoral labor shares, β_{stc} , and value-added shares in 1980. As the first panel of Table 9 shows, we find that the contribution of the shift and share terms to changes in wage-bill shares are very similar in Europe and the US. For example, the share term is 82%, 23% and 25% for low-, middle-, and high-skill occupations in Europe for 73%, 28% and 31%

⁵⁰Due to data availability issues we use 2015 values from the LFS and EU-SILC to compute 2016 factor shares.

⁵¹The average values for the share component are 0.0124, -0.0455 and 0.0429 for low-, middle- and high-skill occupations.

Table 9: Shift-Share Analysis for European Countries

	Wage Bill 1980-2016			Employ. Sh. 1995-2016		
	Low	Middle	High	Low	Middle	High
Total	0.016	-0.18	0.17	0.011	-0.10	0.09
% Shift	18%	77%	75%	-10%	34%	55%
% Share	82%	23%	25%	110%	66%	45%

Notes: Data from EUKLEMS, EU LFS and EU SILC. See discussion in main text and Appendix B.3 for details. Shift and Share report the median contribution in accounting for the total variation.

in the US.⁵²

From the LFS microdata, we obtain data on hours worked by occupation. These data combined with the EUKLEMS allow us to conduct a shift-share analysis of the the distribution of hours worked from 1995 to 2015 in 16 European countries⁵³

$$\frac{X_{jtc}}{X_{tc}} - \frac{X_{j0c}}{X_{0c}} = \sum_{s \in \mathcal{S}} (\rho_{jstc} - \rho_{js0c}) \frac{X_{s0c}}{X_{0c}} + \sum_{s \in \mathcal{S}} \rho_{jstc} \left(\frac{X_{stc}}{X_{tc}} - \frac{X_{s0c}}{X_{0c}} \right), \quad (30)$$

where $\rho_{jstc} \equiv \frac{X_{jstc}}{X_{stc}}$. The second panel in Table 9 reports the median contributions of the shift and share terms across the sample of 16 countries.⁵⁴ We find that the share component accounts for an important fraction of the change in employment shares for all occupations. In particular, it accounts for 110% of the change in the share of hours worked for low-, 66% for middle- and 45% for high-skill occupations. These magnitudes are larger than the contributions of the share term we have computed for the US from 1980 to 2016.

As we have discussed, there are many potential drivers of the sectoral reallocation of value added and hours worked. To directly assess the role of the income-driven channel in Europe, we implement the partial-equilibrium quantification we have conducted for the US in Section 3.⁵⁵ Specifically, we compute the effect of neutral increases in expenditure per capita of the magnitude observed in each country from 1980 to 2016 on the distribution of wage bills across occupations. We calibrate the preference parameters $\{\sigma, \varepsilon_s\}$ to the values we have estimated

⁵²Table A.7 in the Appendix reports the results country by country.

⁵³These countries are: Austria, Belgium, Germany, Denmark, Greece, Spain, Finland, Ireland, Iceland, Italy, Luxembourg, the Netherlands, Norway, Portugal, Sweden, and the UK.

⁵⁴Table A.8 in the Appendix reports the exercise country by country. We note that if we restrict the sample to the 8 countries in our wage bill sample, we find similar patterns.

⁵⁵We decide not to conduct the general equilibrium quantification because of lack of data in the European sample of countries, most prominently on wages by occupation for 1980. Of course, the partial-equilibrium quantification cannot reflect the indirect effect of the income-driven channel on polarization through relative sectoral prices that, as we have shown above, has been significant for the US.

Table 10: Wage Bill Shares for European Countries 1980-2016, Partial Equilibrium Exercise

	Low-skill				Middle-skill				High-high			
	Data	Model Share Term			Data	Model Share Term			Data	Model Share Term		
	Total	All	Only E	$\frac{\text{Only } E}{\text{Data}}$	Total	All	Only E	$\frac{\text{Only } E}{\text{Data}}$	Total	All	Only E	$\frac{\text{Only } E}{\text{Data}}$
Austria	0.013	0.022	0.008	63%	-0.17	-0.04	-0.02	9%	0.16	0.06	0.03	22%
Spain	0.022	0.017	0.012	56%	-0.18	-0.03	-0.02	10%	0.16	0.05	0.04	27%
Finland	0.016	0.033	0.013	78%	-0.19	-0.08	-0.03	17%	0.17	0.12	0.06	33%
France	0.021	0.012	0.006	28%	-0.19	-0.06	-0.02	10%	0.17	0.05	0.03	17%
Italy	0.017	0.015	0.006	34%	-0.17	-0.06	-0.02	10%	0.16	0.04	0.02	14%
Germany	0.014	0.013	0.007	49%	-0.19	-0.03	-0.02	9%	0.18	0.05	0.03	18%
Netherlands	0.010	0.018	0.006	65%	-0.18	-0.04	-0.02	9%	0.17	0.07	0.03	19%
UK	0.023	0.024	0.011	46%	-0.21	-0.06	-0.03	15%	0.19	0.07	0.05	24%
Median EU	0.017	0.017	0.008	52%	-0.19	-0.05	-0.02	10%	0.17	0.05	0.03	20%
US	0.021	0.026	0.012	57%	-0.209	-0.078	-0.025	12%	0.19	0.09	0.05	27%

Notes: The column "Data Total" reports the total change in wage bill shares from 1980 to 2016. Column "All" reports the model-generated Share term when feeding in both changes in relative prices and expenditure. Column " E " reports the model-generated share term when only feeding changes in expenditure.

from the US CEX,⁵⁶ and allow the taste parameters $\{\zeta_s\}$ to vary across countries to match the 1980 value-added shares. We calibrate the factor intensities, $\{\hat{\alpha}_{jstc}, \beta_{tc}\}$, to the values computed above for the shift-share decomposition. We conduct two simulations. In the first, we change both expenditure per capita and relative prices to their observed values from 1980 to 2016. In the second, we only shock the 1980 economy with the increase in expenditure per capita, and keep relative prices at their 1980 levels.

Table 10 reports the change in the share term in these two simulations, and compares them with the actual change in the wage-bill shares in the data. These simulations show that the direct effect of the income-driven channel generates changes in wage-bill shares across occupations that represent a significant part of the changes we have observed in the data. For the median country, the effect of the neutral increase in expenditure in this partial equilibrium setting accounts for 52% of the change in the wage-bill share of the low-skill occupations, 10% of the middle-, and 20% of the high-skill occupations. These magnitudes are very similar to those in the partial equilibrium quantification for the US. Therefore, we conclude that the income-driven channel has played a similar role in the polarization of labor markets in European economies as it has in the US.

⁵⁶Comin et al. (2015) document that the parameters that characterize nonhomothetic CES preferences are stable across countries.

5.2 US Labor-Market Outcomes, 1950-2035

We conclude the paper by analyzing the relevance of our mechanism for the US in a wider time span. We begin looking further backwards into the 1950-1980 period. Then, we project what our mechanism entails going forward into the next fifteen years. For these exercises, we leverage on the property documented in [Comin et al. \(2015\)](#) that the demand parameters $\{\sigma, \varepsilon_s\}_{s \in \mathcal{S}}$ of nonhomothetic CES are stable for very different income levels, and keep them fixed at their baseline values.⁵⁷

5.2.1 Back-tracking 1950-1980

Much of the job polarization debate has focused on the post-1980 period when the automation of production processes and globalization forces have been at their apex. Yet, [Bárány and Siegel \(2018\)](#) suggest that labor-market polarization may have well started in the US by the 1950s.⁵⁸ A natural question to ask is what drove the US evolution of wages and hours worked across occupations during the 1950-1980 period which did not witness such strong biased technical change and globalization. In particular, what role did the income-driven channel play?

We explore this question with the calibrated general equilibrium used in [Section 4.3](#). We follow the same approach to calibrate the model as described in [Section 4.2](#) with 1950 and 1980 as initial and final years. We keep the preference parameters to the values estimated from [Table 3](#). The taste parameters $\{\zeta_s\}_{s \in \mathcal{S}}$ are set to match the 1950 sectoral value-added shares. The factor intensities and labor shares are set to match those observed in the data. We calibrate the E -shock so that the model's level of aggregate expenditure per capita and the Fisher price index match the levels observed in 1980.

[Tables A.9 and A.10](#) in [Appendix A](#) present the results from our simulation. The main conclusion is that the income-driven channel contributed substantially to the changes in labor-market outcomes from 1950 to 1980. In particular, it caused an increase in the wage-bill share of high-skill occupations of 4.5 percentage points, which represents around 40% of the actual increase observed in the data. This contribution reflects both the increase in the share of hours worked (by 2.3 percentage points) and in the wage of high- relative to middle-skill occupations from 1.15 in 1950 to 1.22 in 1980. The income-driven channel also

⁵⁷[Comin et al. \(2015\)](#) show that the estimated values of $\{\sigma, \varepsilon_s\}_{s \in \mathcal{S}}$ are stable across different income levels both using micro and aggregate data. In particular, they estimate very similar nonhomotheticities parameters $\{\varepsilon_s\}_{s \in \mathcal{S}}$ using the Indian National Sample Survey and the US Consumer Expenditure Survey.

⁵⁸Our data are consistent with this observation for relative wages, while for employment we observe a slight decline in employment shares of low-skill occupations over the 1950-1980 period. As pointed out by [Cerina et al. \(2017\)](#), this discrepancy is due to the fact that we include all occupations in our sample, while [Bárány and Siegel](#) exclude agricultural occupations.

led to an increase in the wage-bill share of low-skill occupations by 2.5 percentage points, both through an increase in the share of hours worked (2.4 percentage points) and in the relative wage which increased from 70% of the wage accrued by middle-skill occupations to 75%.⁵⁹ Finally, the income-driven channel also contributed to the decline of the share of hours worked by middle-skill occupations (4.7 percentage points of a total of 7.8), but did not contribute to the decline in the wage-bill share accrued by middle-skill occupations.

5.2.2 Looking into the future of US labor markets: 2016-2035

We conclude by using our framework to forecast labor-market outcomes in 2035. We calibrate our model to 2016 using an empirical strategy analogous to our baseline calibration for 1980. As a result, our calibrated model matches the wage-bill shares, employment shares and relative wages in 2016.⁶⁰ We then simulate an E shock so that household expenditures and the Fisher price index increase by half of the observed increase from 1980 to 2016. While it seems reasonable to assume that household expenditures will increase over the next fifteen years, it is less clear what patterns innovation, offshoring and other elements of the biased-technology shock will follow. For this reason, we keep constant factor intensities, labor shares and relative sectoral TFP at their 2016 levels.

Table A.11 in the appendix shows that our model predicts an exacerbation of labor-market polarization. The model suggests that the wage-bill shares of low- and high-skill occupations will increase from 8.8% to 9.9% and from 49% to 53%, while the wage-bill share of middle-skill occupations will decrease from 42% to 37%. This polarization of the distribution of wage bills will be reflected in both a polarization of hours worked and wage rates across occupations. The share of hours worked by low- and high-skill occupations will increase from 12.9% to 14.2% and from 38% to 41%. Similarly, the average wage earned in low- and high-skill occupations relative to the average wage earned by middle-skill occupations will increase from 0.8 to 0.84, and from 1.49 to 1.57.

Comparing these changes with those produced by the income-driven channel from 1980 to 2016, we observe that the forecasted polarization increases more than proportionally relative to the size assumed for the E shock. This convexity of polarization in the E shock arises because, as the economy becomes richer, household expenditures concentrate further on the most expenditure-elastic sectors. This result is connected to our discussion of the expenditure elasticity formula in Equation (16). There, we showed that, as total household

⁵⁹In the data, we observe a slight decline in the labor share and wage-bill shares of low-skill occupations. This is due mostly to the shift term, which reflects the decline in the factor intensity of low-skill occupations across sectors.

⁶⁰We obtain very similar predictions if we keep running forward our baseline model with the same shock rather than recalibrating it to match 2016 exactly.

expenditure rises, some sectors which used to be luxuries become necessities (i.e., the expenditure elasticity goes from above to below 1). Moreover, goods that were already necessities reduce their expenditure elasticities even further. This induces a non-linear effect of income on expenditure shares whereby expenditure shares only grow in the remaining luxury sectors and decrease in all other sectors. This is indeed what we observe in our calibrated model. The expenditure elasticity of the Government sector goes below 1, and it further decreases for sectors that were already a necessity, such as Agriculture or Manufacturing. Only the three most expenditure-elastic sectors remain with an expenditure elasticity above 1 by year 2035. Since these are precisely the sectors that are more intensive in low- and high-skill occupations, the effect of the E shock is exacerbated relative to our baseline period.

6 Conclusion

This paper has explored the role of the income-driven channel in labor-market polarization. We have developed a general-equilibrium model that allows for heterogeneity in expenditure elasticities of demand across sectors, endogenous job assignment, time-varying factor intensities, and sector-specific TFP. This last two elements of our model capture standard drivers of polarization such as factor-biased technical change, routinization and other capital-skill complementarities, and offshoring. Importantly, we have calibrated the key preference parameters using micro-level estimates and, hence, we have pinned them down independently of the aggregate phenomena we have explored with the model.

Our quantitative exercise has revealed that the income-driven channel has played a significant role in the polarization of US labor markets. It has been the main driver of the increase in the share of low-skill occupations in the total wage bill and total hours worked. Specifically, the income-driven channel has caused 90% and 64% of the increase in the share of low-skill occupations in total wage bill and hours worked, 2% and 34% of the decline in the share of middle-skill occupations in total wage bill and hours worked, and 35% and 28% of the increase in the share of high-skill occupations in total wage bill and hours worked. Additionally, the income-driven channel accounts for 46% and 29% of the increases of the average wage of low- and high-skill occupations relative to middle-skill occupations from 1980 through 2016.

We have extended our analysis to explore the drivers of polarization in the US during the 1950-80 period, in European countries during the 1980-2016 period, and to forecast the evolution of polarization in the US over the next fifteen years. In the first two exercises, we find that the income-driven channel has played remarkably similar roles for polarization to the US over the 1980-2016 period. Our simulations for the next fifteen years suggest that

polarization coming from the income-driven channel will continue at least at the same pace as in the past forty years.

Finally, this paper has taken the evolution of factor intensities and sectoral TFP as exogenous, and focused on identifying the effect of nonhomotheticities on polarization. However, the direction of innovation is endogenous—as the literature on endogenous technological change points out. This dichotomy between technological progress and income-driven polarization is, at least in part, fallacious. The direction of innovation is endogenous and responds to the size of the market. Therefore, in a nonhomothetic world, sectoral TFP and factor intensities depend on the level of income. [Comin et al. \(2016\)](#) have undertaken part of this analysis in the context of a growth model. Extending their framework to heterogeneous occupations may provide new avenues for the income-driven channel to affect labor-market polarization.

References

- ACEMOGLU, D. (1999): “Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence,” *American Economic Review*, 89, 1259–1278.
- ACEMOGLU, D. AND D. AUTOR (2011): *Skills, Tasks and Technologies: Implications for Employment and Earnings*, Elsevier, vol. 4 of *Handbook of Labor Economics*, chap. 12, 1043–1171.
- AGUIAR, M. AND M. BILS (2015): “Has Consumption Inequality Mirrored Income Inequality?” *American Economic Review*, 105, 2725–56.
- AUTOR, D. AND D. DORN (2009): “This Job Is “Getting Old”: Measuring Changes in Job Opportunities Using Occupational Age Structure,” *American Economic Review*, 99, 45–51.
- AUTOR, D. H. AND D. DORN (2013): “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 103, 1553–97.
- AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): “The Polarization of the U.S. Labor Market,” *American Economic Review*, 96, 189–194.
- (2008): “Trends in U.S. Wage Inequality: Revising the Revisionists,” *The Review of Economics and Statistics*, 90, 300–323.
- AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, 118, 1279–1333.
- BÁRÁNY, Z. AND C. SIEGEL (2019): “Engines of Sectoral Labor Productivity Growth,” 2019 Meeting Papers 872, Society for Economic Dynamics.

- BÁRÁNY, Z. L. AND C. SIEGEL (2018): “Job Polarization and Structural Change,” *American Economic Journal: Macroeconomics*, 10, 57–89.
- BASCO, S. AND M. MESTIERI (2013): “Heterogeneous trade costs and wage inequality: A model of two globalizations,” *Journal of International Economics*, 89, 393–406.
- BUERA, F. J. AND J. P. KABOSKI (2012): “The Rise of the Service Economy,” *American Economic Review*, 102, 2540–69.
- BUERA, F. J., J. P. KABOSKI, AND R. ROGERSON (2015): “Skill Biased Structural Change,” Working Paper 21165, National Bureau of Economic Research.
- BUERA, F. J., J. P. KABOSKI, R. ROGERSON, AND J. VIZCAINO (2018): “Skill Biased Structural Change,” Working paper, CEPR.
- CARON, J., T. FALLY, AND J. MARKUSEN (2020): “Per capita income and the demand for skills,” *Journal of International Economics*, 123.
- CASELLI, F. (1999): “Technological Revolutions,” *American Economic Review*, 89, 78–102.
- CERINA, F., A. MORO, AND M. PETERSEN RENDALL (2017): “The role of gender in employment polarization,” LSE Online Documents on Economics 86170, LSE.
- COMIN, D., D. LASHKARI, AND M. MESTIERI (2016): “Demand-Pull, Technology-Push, and the Sectoral Direction of Innovation,” 2016 Meeting Papers 1287, Society for Economic Dynamics.
- COMIN, D. A., D. LASHKARI, AND M. MESTIERI (2015): “Structural Change with Long-run Income and Price Effects,” NBER Working Papers 21595, National Bureau of Economic Research. Accepted at *Econometrica*.
- CORTES, G. M., N. JAIMOVICH, AND H. E. SIU (2017): “Disappearing routine jobs: Who, how, and why?” *Journal of Monetary Economics*, 91, 69–87.
- COZZI, G. AND G. IMPULLITTI (2016): “Globalization and Wage Polarization,” *The Review of Economics and Statistics*, 98, 984–1000.
- CRAVINO, J. AND S. SOTELO (2017): “Trade Induced Structural Change and the Skill Premium,” NBER Working Papers 23503, National Bureau of Economic Research, Inc.
- FRENCH, K. (2008): “Presidential Address: The Cost of Active Investing,” *Journal of Finance*, 63, 1537–1573.
- GALLE, S., A. RODRÍGUEZ-CLARE, AND M. YI (2017): “Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade,” NBER Working Papers 23737, NBER, Inc.
- GOOS, M. AND A. MANNING (2007): “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain,” *The Review of Economics and Statistics*, 89, 118–133.
- GOOS, M., A. MANNING, AND A. SALOMONS (2009): “Job Polarization in Europe,” *American Economic Review*, 99, 58–63.
- (2014): “Explaining Job Polarization: Routine-Biased Technological Change and

- Offshoring,” *American Economic Review*, 104, 2509–2526.
- HANOCH, G. (1975): “Production and Demand Models with Direct or Indirect Implicit Additivity,” *Econometrica*, 43, 395–419.
- HERRENDORF, B., R. ROGERSON, AND A. VALENTINYI (2013): “Two Perspectives on Preferences and Structural Transformation,” *American Economic Review*, 103, 2752–2789.
- HSIEH, C.-T., E. HURST, C. I. JONES, AND P. J. KLENOW (2013): “The Allocation of Talent and U.S. Economic Growth,” NBER Working Papers 18693, NBER, Inc.
- IKENAGA, T. AND R. KAMBAYASHI (2016): “Task Polarization in the Japanese Labor Market: Evidence of a Long-Term Trend,” *Industrial Relations: A Journal of Economy and Society*, 55, 267–293.
- JAIMOVICH, N., S. REBELO, AND A. WONG (2019): “Trading down and the business cycle,” *Journal of Monetary Economics*, 102, 96 – 121, carnegie-Rochester-NYU Conference Series.
- KATZ, L. F. AND K. M. MURPHY (1992): “Changes in Relative Wages, 1963-1987: Supply and Demand Factors,” *The Quarterly Journal of Economics*, 107, 35–78.
- KRUEGER, D. AND F. PERRI (2006): “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory1,” *The Review of Economic Studies*, 73, 163–193.
- KRUSELL, P., L. E. OHANIAN, J.-V. RÍOS-RULL, AND G. L. VIOLANTE (2000): “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68, 1029–1053.
- LEE, S. Y. AND Y. SHIN (2017): “Horizontal and Vertical Polarization: Task-Specific Technological Change in a Multi-Sector Economy,” Working Paper 23283, National Bureau of Economic Research.
- LEONARDI, M. (2015): “The Effect of Product Demand on Inequality: Evidence from the United States and the United Kingdom,” *American Economic Journal: Applied Economics*, 7, 221–47.
- LIND, N. AND N. RAMONDO (2018): “Trade with Correlation,” NBER Working Papers 24380, National Bureau of Economic Research, Inc.
- MANNING, A. (2004): “We Can Work It Out: the Impact of Technological Change on the Demand for Low Skill Workers,” Cep discussion papers, CEP, LSE.
- MAZZOLARI, F. AND G. RAGUSA (2013): “Spillovers from High-Skill Consumption to Low-Skill Labor Markets,” *The Review of Economics and Statistics*, 95, 74–86.
- NGAI, L. R. AND O. SEVINC (2020): “A Multisector Perspective on Wage Stagnation,” Working Paper 14855, CEPR.
- SCHIMMELPFENNIG, A. (1998): “Skill-biased technical change vs. structural change: Insights from a new view of the structure of an economy,” Kiel Institute Working Papers 868.