The Income Channel to Labor-Market Polarization*

Diego Comin
Dartmouth College

Ana Danieli
Amazon

Marti Mestieri
IAE-CSIC

May 31, 2026

Abstract

We study labor-market polarization by developing a framework with heterogeneous consumers, non-homothetic preferences, and endogenous labor supply à la Roy (1951). Because income-elastic sectors are intensive in high- and low-skilled occupations, shocks that raise either the average or the dispersion of income across households shift relative labor demand toward these occupations, generating an income channel to polarization. We quantify this channel by comparing the effect of shocks that affect the distribution of household income in our model and in a version with homothetic preferences in which the income distribution does not affect sectoral composition. The income channel is sizable: it explains 42% and 210% of the increases in the relative wages of high- and low-skill workers observed between 1980 and 2016, and 26% and 51% of the increase in their hours-worked shares. A quarter of its effect operates through income dispersion, highlighting the importance of departing from representative-consumer frameworks in non-homothetic environments.

Keywords: Labor-market polarization, non-homothetic demand, heterogeneous households.
JEL Classification: E21, E23, J23, J31.

*A previous version of this paper was circulated with the title "Income-driven Labor-Market Polarization." We thank Gadi Barlevy, David Dorn, Matthias Doepke, Thibault Fally, Koichiro Ito, Ryan Kellogg, Kiminori Matsuyama, Rachel Ngai, Juan Pablo Nicolini, Richard Rogerson, Yongs Shin, Michael Sposi, Christian Siegel, Chris Tonetti, Joseph Zeira, Nathan Zorzi, and attendees at various seminars and conferences for helpful comments and discussions. We are grateful to Lara Pons for her outstanding research assistance. Comin expresses his gratitude for the financial support from the European Commission under the FRAME project and from the NSF. Mestieri acknowledges financial support from the Europa Excelencia and the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics CEX2024-001476-S) funded by the MCIN/AEI. All errors are our own.

1 Introduction

Labor markets in the U.S. and Western Europe have become increasingly polarized since the 1980s. The wages of high- and low-skilled occupations have increased relative to those of middle-skilled occupations. Simultaneously, the share of hours worked in middle-skilled occupations has decreased while the shares of both high- and low-skilled occupations have increased. Consequently, the hollowing out of middle-skilled occupations has reduced their share in the total wage bill in the U.S. from 63% in 1980 to 41% in 2016.

The social significance of these trends has given rise to a vast literature that has emphasized two main drivers of labor-market polarization: technological change and international offshoring. A large body of work documents polarization and shows that the adoption of computers and other digital technologies increased the productivity of high-skilled workers while automating routine tasks that were typically conducted by middle-skilled labor (see Acemoglu, 1999; Autor et al., 2003; Goos and Manning, 2007; Goos et al., 2009, 2014; Autor et al., 2006, 2008; Autor and Dorn, 2009, 2013, among others). A complementary line of research suggests that globalization can reshape labor demand by changing the tradability of tasks. Globalization contributes to labor-market polarization because declining offshoring costs disproportionately affect middle-skill-intensive tasks (see Feenstra and Hanson, 1999; Grossman and Rossi-Hansberg, 2008; Blinder, 2009; Jensen and Kletzer, 2010; Firpo et al., 2011; Blinder and Krueger, 2013, among others).[1]

In this paper, we propose and quantify a new channel to explain labor-market polarization. Through our mechanism, changes in the distribution of income across households, such as shifts in the mean or the dispersion, can increase the polarization of labor-market outcomes. We refer to this mechanism as the income channel of polarization.

The income channel builds on a novel empirical fact: expenditure-elastic sectors are intensive in high- and low-skill occupations in production, whereas there is no such correlation for middle-skill occupations. Because preferences are non-homothetic, any shock that raises average income or increases income dispersion, shifting income from low- to high-income households, induces a reallocation of demand from sectors with low expenditure elasticity to sectors with high expenditure elasticity. Since these high-expenditure-elasticity 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. The result is a hollowing out of middle-skill occupations and the polarization of workers' earnings.

[1] Technology and globalization are connected. For example, Basco and Mestieri (2013) document that trade costs have fallen more in middle-skill-intensive industries, which are also more intensive in information technologies and routine tasks. They show that this can generate wage polarization in rich countries. See Cozzi and Impullitti (2016) for a dynamic account of polarization and trade in a growth model.

To study the drivers of labor market polarization and the role of the income channel, we develop a multi-sector general equilibrium framework with three key ingredients: (i) heterogeneous consumers, (ii) non-homothetic preferences as in Hanoch (1975) and Comin et al. (2021), and (iii) endogenous labor supply à la Roy (1951), augmented with a Pareto mixture to replicate the top-income share.

This framework allows us to capture three channels through which shocks may affect labor-market polarization. First, as in the models of labor market polarization cited above, certain shocks, such as changes in the intensity of occupations within sectors or in the dispersion of productivity draws across occupations, directly affect labor-market outcomes by shifting relative labor demand or supply across occupations.

Second, as in models of structural transformation and labor-market polarization (e.g., Lee and Shin, 2017 and Bárány and Siegel, 2018), the framework incorporates indirect effects of shocks that operate through changes in the sectoral composition of the economy. Shocks to relative sectoral TFP or to taste parameters alter the sectoral composition of the economy and, because factor intensity varies across sectors, these changes shift relative labor demand across occupations.[2]

Third, shocks may affect labor market outcomes by altering the distribution of income across households. In particular, changes in the dispersion of productivity draws across occupations, aggregate TFP and factor intensities of occupations affect both the average and the dispersion of income across households. With non-homothetic preferences, these changes in the distribution of income induce shifts in the sectoral demand and, in turn, in the relative labor demand across occupations.

We discipline the model using micro-level preference estimates from the CEX and calibrate the remaining parameters to match key moments of U.S. labor markets in 1980 and 2016—specifically, relative wages, hours-worked shares by occupation, sectoral value-added distribution, variance of log labor income, and top-income share.

We first use this quantitative framework to study the contribution of specific shocks—calibrated changes in the model's economic primitives between 1980 and 2016, which we refer to interchangeably as impulses—to the changes in relative wages, hours worked and wage-bill shares across occupations between 1980 and 2016. Changes in sectoral factor intensities explain all the increase in the relative wage of high-skilled workers (vs. middle-skilled workers) and low-skilled workers, and 70% and 40% of the increases in their hours-worked shares,

[2]International trade is another force that can affect domestic sectoral composition, especially in tradable sectors. In a previous version of the paper, Comin et al. (2020), we introduced trade wedges based on sectoral net exports to relax the equality between domestic absorption and domestic production. The main quantitative conclusions about the income channel were unchanged. We therefore abstract from trade in the present analysis.

respectively. By contrast, changes in the variance of productivity draws across occupations, the Pareto tail, and aggregate TFP do not affect the relative wages of high-skilled workers and reduce the relative wages of low-skilled workers. However, they account for 31% and 91% of the increases in their hours-worked shares, respectively. Relative sectoral TFP explains a small share of the increases in relative wages and hours worked of high-skilled workers (13% and 7.6%), but it accounts for substantial shares of the increases for low-skilled workers (101% and 27.5%).

However, studying the effects of impulses separately does not reveal the contributions of the three channels to labor-market polarization, because several impulses (e.g., factor intensity, the dispersion of productivity draws, and aggregate TFP) simultaneously affect both the direct channel and the income channel.

To assess the role of the income channel, we develop two alternative versions of the model that progressively shut it down. One version replaces heterogeneous consumers with a representative consumer with non-homothetic preferences, thereby restricting the income channel to operate only through changes in the economy's average household income. The other replaces the non-homothetic CES preferences with standard homothetic CES preferences, eliminating the income channel entirely since changes in the distribution of income no longer affect the sectoral composition of the economy. Finally, we show in a lemma that the calibration strategy is block-recursive, implying that all shocks and parameters affecting the direct and income channels are identical across the different versions of the model.

We quantify each of the three channels in our framework. The direct channel is measured by the effect of applying the four shocks that directly influence labor supply and demand—changes in factor intensity, the dispersion of productivity draws, the Pareto tail, and aggregate TFP—in the homothetic model. The income channel is measured by the differential effect of the shocks that alter the distribution of income across households—changes in factor intensity, the dispersion of productivity draws, the Pareto tail, and aggregate TFP—in the baseline model relative to the homothetic model. Finally, the sectoral-composition channel is measured by the effect of the shocks that modify the sectoral composition of the economy (but do not directly affect the income distribution)—relative sectoral TFP and taste parameters—in the baseline model.

The income channel is a major driver of the rise in relative wages for high- and low-skilled workers. It contributes 8 log points to the increase in each group's relative wage, corresponding to 42% and 210% of the observed increases, respectively. The 210% exceeds 100% because the direct channel and the sectoral-composition channel push the low-skilled relative wage downward, partially offsetting the upward contribution from the income channel. The direct channel contributes 13 log points for high-skilled relative wages and -1.4 log points

for low-skilled relative wages, while the sectoral composition channel reduces both relative wages.

The income channel also plays a significant role in driving the reallocation of labor toward high- and low-skilled occupations. It raises the high-skilled hours-worked share by 3 percentage points (p.p.) and the low-skilled share by 1.4 p.p., accounting for 26% and 51% of the actual increases, respectively. In contrast, the direct channel explains 79% and 71% of the increases in the hours-worked shares of high- and low-skilled workers, while the sectoral-composition channel reduces both shares.

We study the importance of heterogeneous consumers for the income channel by comparing its magnitude in the baseline and in the version with a representative consumer with non-homothetic preferences. The contribution of the income channel is cut by approximately one quarter when we omit the effects on polarization that operate through the dispersion of household income.

The role of consumer heterogeneity in amplifying shocks has previously been emphasized in models featuring liquidity constraints and heterogeneous marginal propensities to consume (e.g., Werning, 2015, Auclert, 2019, and Patterson, 2023). Our framework provides a new rationale for the importance of heterogeneity, rooted in non-homothetic preferences: consumption functions for income-elastic sectors are convex in household expenditure, leading to the amplification of shocks that increase the variance of expenditure across households. These findings underscore that the algebraic convenience of representative-consumer models may come at the cost of substantially understating the amplification produced by heterogeneous consumer behavior.

Related literature In addition to the papers already cited, our paper is related to three strands of the literature. First, the so-called consumption externality or trickle-down hypothesis formulated by Manning (2004) argues that the demand for low-skilled workers is tied to the income of high-skilled ones. When the incomes of affluent households rise, these households substitute home-produced services with their market counterparts. Since home-produced services are intensive in low-skill workers, this substitution increases the demand for low-skill market services, boosting the wages and employment of low-skill workers.[3] This consumption externality is a manifestation of the income channel, but it captures only one component of it.

Several studies, including Manning (2004), Mazzolari and Ragusa (2013), Leonardi (2015) and Autor and Dorn (2013), have assessed the consumption externality by estimating reduced-form regressions, at the local labor market level, of changes in low-skill employment or wages

[3]See Murphy (2016) for an example where they move in opposite directions.

on the income or wage bill of top-income-decile households. The results from these exercises are mixed. More importantly, this empirical strategy does not properly capture the income channel for two reasons. First, it focuses narrowly on the top decile, whereas the income channel is activated by any change in the distribution of income across households (e.g., changes in the mean or dispersion). Second, it misses shocks that are common across local labor markets, such as nationwide increases in average income or inequality, which also shift sectoral demand and laborâĂŚmarket outcomes.

In contrast, we use cross-sectional household variation to estimate expenditure elasticities for different sectors. We use these estimates to calibrate our model and run counterfactuals for the combination of shocks that drive the income distribution. In particular, we compare outcomes in the baseline with non-homothetic preferences, where the income channel is present, with those in the homothetic version in which the income channel is shut down. We find that the income channel is quantitatively important for both the relative wages and hours-worked shares of high- and low-skill occupations.

Second, our paper is related to studies that have linked structural change and inequality. Lee and Shin (2017), Bárány and Siegel (2018, 2019) and Ngai and Sevinc (2020) have built multi-sector general equilibrium models that incorporate both the direct and the sectoral composition channels to labor-market polarization. However, because they have homothetic preferences, they abstract from the income channel.

Few papers have introduced non-homotheticities in multi-sector models of labor market polarization providing a role for the income channel. One exception is Buera et al. (2021), who introduce Stone-Geary preferences in a representative agent model with two types of skills and exogenous labor supply. There are, however, important methodological differences with our paper, including elements of the model (preferences, endogenous labor supply, number of occupations, sectors, and non-representative consumer) calibration strategy, focus (skill premium vs. labor market polarization) and the measurement of the income channel. As a result, we find a significantly larger role for the income channel in the increase of the relative wage of skilled workers (i.e., skill premium).[4]

Third, our paper is also related to a small group of papers that have incorporated non-

[4]Cravino and Sotelo (2019) study the effect of trade on the skill premium in a model with two sectors, two skills, a representative consumer, exogenous labor supply and non-homothetic CES preferences. Though this setting is closer to ours, there are still important differences, endogenous labor supply à la Roy, heterogeneous consumers within occupations, and the focus is quite different as we are interested in quantifying the different channels to labor market polarization and Cravino and Sotelo (2019) are interested in quantifying the role of globalization for the skill premium. Caron et al. (2020) develop a multi-country, multi-sector model with trade with non-homothetic demand and different intensities of skilled and unskilled workers across sectors. They study the effect of sectoral TFP growth and reduction of trade barriers on the skill premium. This paper presents significant differences with ours in the setting and analysis including that they consider only traded sectors. They find a small effect ($<1\%$) on the skill premium in the US.

homothetic preferences in heterogeneous-agent models to obtain greater amplification of business cycle shocks.[5] Jaimovich et al. (2019) introduce a model with non-homotheticities in product quality. In each sector there are high- and low-quality products, and low-quality products are less labor intensive. As a result, in recessions, consumers shift their demand towards low-quality products, further depressing labor demand. Andreolli et al. (2024) develop a model with two sectors (essential and non-essential) that differ in how intensively they employ two types of agents with different marginal propensity to consume (hand-to-mouth vs. permanent-income/Euler agents). They then study the model's consistency with a new fact: the income of non-essential workers responds more than that of essential workers to monetary-policy shocks. Our paper differs from these in two main respects. First, the models differ in key details: preferences, the number of sectors and occupations, formulation of labor supply and presence of within-occupation heterogeneity in income and consumption. Second, the focus differs: these papers study how non-homotheticities amplify specific business-cycle shocks, whereas we quantify the role of the channels in labor market polarization.

The rest of the paper is organized as follows. Section 2 documents labor-market polarization and the key empirical facts that motivate the income channel. Section 3 develops the model. Section 4 presents the calibration strategy. Section 5 contains the main quantitative analysis: it assesses the effects of individual impulses on polarization, develops representative-consumer and homothetic versions of the model, quantifies the three channels behind labor-market polarization, and discusses policies that affect sectoral composition and the income channel. Section 6 concludes.

2 Facts about the polarization of labor markets

This section documents the key facts for our study. We begin with definitions and details on data construction. We then use the data to (i) document the polarization of labor markets and (ii) present the empirical relationships that motivate the channels used by the model to account for labor-market polarization.

2.1 Data Construction

We follow Acemoglu and Autor (2011) in constructing our data on wages and hours-worked shares by occupation. Hours-worked data come from IPUMS USA and include the decennial censuses from 1980 to 2000 (in 10-year intervals) and annual data from the American Community Survey (ACS) from 2000 to 2016. The sample is restricted to individuals aged 16-64

[5] Additionally, Orchard (2025) introduces non-homothetic preferences to study the cyclicality of sectoral inflation, and Olivi et al. (2025) study how non-homotheticities affect optimal monetary policy.

who were employed in the previous year and assigned to a known occupation (i.e., occupation is not missing or unemployed). We further exclude the top and bottom 5% of the hourly wage distribution.

Wage data come from the Current Population Survey (CPS). Occupations and industries are coded using the 1990 Census Bureau classification scheme, which provides a consistent taxonomy across all sample years. Following Acemoglu and Autor (2011), the 382 original occupations are grouped into four broader categories based on their average hourly wage in 1980, which serves as a proxy for skill level: (1) managerial, professional, and technical occupations; (2) sales, clerical, and administrative support occupations; (3) production, craft, repair, and operative occupations; and (4) service occupations. High-skill occupations correspond to group 1; middle-skill, to groups 2 and 3; and low-skill, to group 4.

Industries are mapped to the BEA industry classification through their mutual correspondence with NAICS codes. We then aggregate them into eight broad sectors: (i) health and education, (ii) arts, recreation, entertainment, and food services, (iii) finance, information, professional services, and other services, (iv) government, (v) manufacturing, (vi) retail, wholesale, and transportation, (vii) construction, and (viii) agriculture.[6]

We compute the average hourly wage for each of the three skill groups (high-, middle-, and low-skill), W_{jt}.[7] We also compute the total number of hours worked in each occupation-sector pair, L_{jst}. Thus, throughout the analysis the quantity object is total hours worked, not headcount employment. The wage bill for each occupation and sector, WB_{jst}, is then defined as the product of the average wage and the total hours worked. Adding across sectors, we construct occupation-level hours worked, L_{jt}, and wage bill levels, WB_{jt}.

2.2 Labor-market polarization

We use these variables to document the polarization of labor markets in the United States between 1980 and 2016. The wage of high-skilled occupations relative to middle-skilled occupations increased by 19.5 log points, from 1.26 in 1980 to 1.53 in 2016. For low-skilled

[6] As we explain in further detail in Appendix A, we map detailed industries into the BEA Input-Output (I-O) classification using their correspondence with NAICS codes, and then aggregate to eight sectors as follows. Health and Education corresponds to BEA I-O sector 6 (education and health care), plus state and local government value added in educational services and hospitals/health services (BEA U. Value Added by Industry, lines 185-186). Arts, Entertainment, Recreation, and Food Services corresponds to BEA I-O sector 7. Finance, Professional, Information and Other Services aggregates BEA I-O groupings FIRE, PROF, 51, and 81. Government corresponds to BEA I-O sector G, excluding the education and health components reassigned to Health and Education. Manufacturing, Mining and Utilities aggregates BEA I-O sectors 21, 22, and 31G. Retail, Wholesale Trade and Transportation aggregates BEA I-O sectors 42, 44RT, and 48T. Construction corresponds to BEA I-O sector 23, and we assign real estate (line 129) to this sector. Agriculture corresponds to BEA I-O sector 11.

[7] All calculations use the statistical sampling weights.

Table 1: Labor Market Polarization in the United States, 1980 and 2016

	1980	2016	Diff.
Relative Wages			Δ log
High/Middle	1.26	1.53	19.5
Low/Middle	0.77	0.80	3.9
Hours-Worked Shares (%)			Δ Sh.
High	25.2	38.3	13.1
Middle	65.3	48.8	-16.5
Low	9.5	12.9	3.4
Wage-Bill Shares (%)			Δ Sh.
High	30.5	49.9	19.4
Middle	62.5	41.4	-21.1
Low	7.0	8.7	1.8

Notes: This table summarizes the labor-market polarization facts described in the text. Employment shares are shares of total hours worked, wage-bill shares are shares of total compensation, and relative wages are average hourly wages relative to middle-skill occupations. Employment data come from IPUMS USA and the American Community Survey; wage data come from the Current Population Survey.

occupations, the wage relative to middle-skilled occupations increased by 3.9 log points over the same period, from 0.77 to 0.80 (see Table 1).

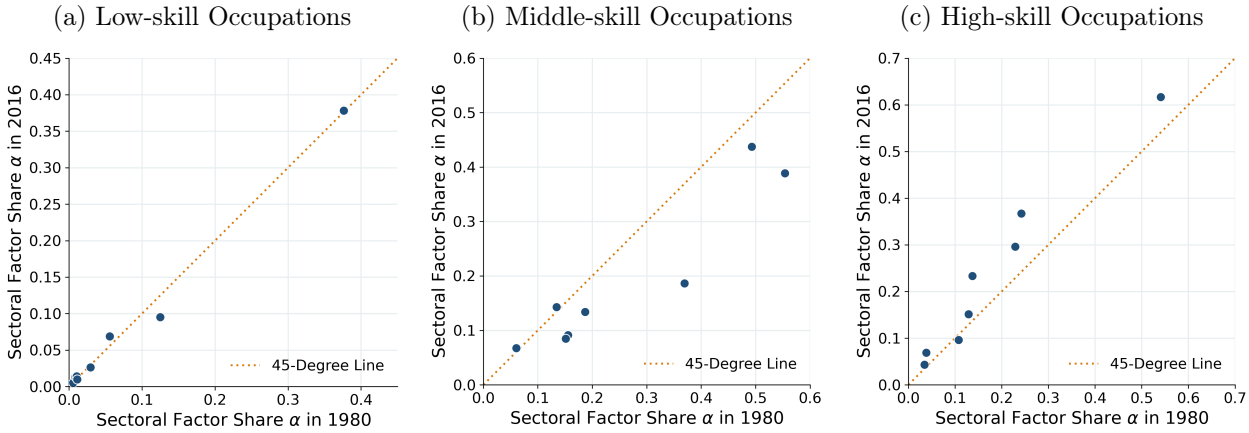
Polarization is also evident in the changing distribution of hours worked. High-skilled occupations expanded their share of total hours by 13 p.p., from 25% in 1980 to 38% in 2016, whereas low-skilled occupations increased their share by 3.4 p.p., from 9.5% to 12.9% over the same period.

The polarization of wages and hours has, in turn, produced a pronounced polarization in the wage bill. The high-skilled wage-bill share rose by 19.4 p.p., from 30% in 1980 to 50% in 2016, while the low-skilled share saw a modest rise of 1.8 p.p., from 7% to 9%. In contrast, the middle-skilled wage-bill share fell sharply by 21.1 p.p., from 63% to 41%. These patterns of polarization are not unique to the US. Similar processes of labor-market polarization have also been documented in Western European economies (Comin et al., 2020).

2.3 Changes in sectoral factor intensity

Polarization arises when demand for high- and low-skilled occupations increases. Changes in relative demand may result from increases in the intensity of sectoral production in high- and low-skilled workers. Figure 1 shows the intensity of each occupation in each sector, measured

Figure 1: Sectoral Factor Shares α in 1980 and 2016



Notes: The sectoral factor share α_{jst} of occupation j in sector s at time t is the wage bill of occupation j in sector s over the total wage bill in sector s. Labor input is computed using hours worked from the Census/ACS, and wages are from the CPS (see Appendix A).

by its wage-bill share, in 1980 and 2016. For high-skilled occupations, the markers tend to lie above the 45-degree line in most sectors, with the largest increases in professional services, manufacturing, and construction.

The increase in the intensity of high-skilled occupations is consistent with biased technological change and offshoring, as both tend to favor these occupations. It is also consistent with non-homothetic production structures, in which the relative importance of different occupations evolves with a firm's scale, as in Comin et al. (2026).

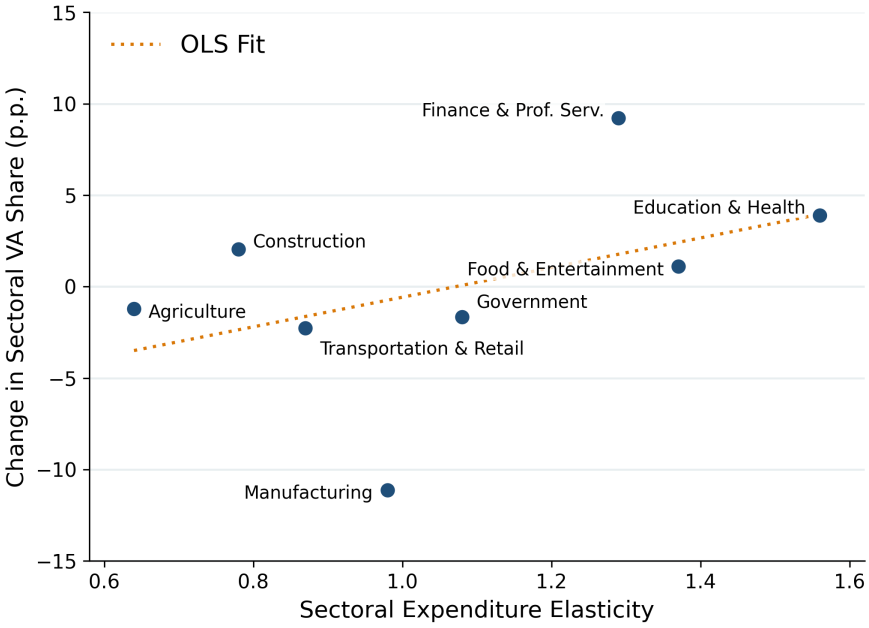
For middle-skilled occupations, the markers tend to lie below the 45-degree line, indicating a reduction in their intensity in production. For low-skilled occupations, the markers tend to lie closer to the 45-degree line, suggesting that changes in their sectoral intensity are not the main driver of the rising demand for low-skilled occupations.

It is also worth noting that, despite the changes in within-sector factor intensity, the rankings of industries by their intensity in any given occupation are the same in 1980 and in 2016. As a result, the associations we present between expenditure elasticities and factor intensity in Section 2.5 are robust to the year used for sectoral factor intensities.

2.4 Expenditure elasticities across sectors

We next provide evidence on the key ingredients underlying the income channel of labor-market polarization. A central component is the sectoral heterogeneity in the expenditure elasticities of demand. To estimate these elasticities, we follow the method developed in Aguiar and Bils (2015), defining expenditures over value-added consumption as in Buera et al.

Figure 2: Change in Sectoral Value-Added Shares 1980-2016 and Expenditure Elasticity



Notes: Sectoral value added measures are from the BEA. Expenditure elasticities are reported in column (4) in Table 3 and are estimated from cross-sectional household data using the Aguiar and Bils (2015) methodology (see Section 3.2).

(2015) and Comin et al. (2021).[8] The estimated expenditure elasticities vary substantially across the eight sectors, ranging from 0.63 in agriculture to 1.54 in health and education.

Comin et al. (2021) argue that income effects are the main driver of sectoral growth, and that expenditure elasticities therefore explain most of the cross-sectoral variation in growth. Figure 2 is consistent with this view. The correlation between a sector's expenditure elasticity and its change in the value-added share between 1980 and 2016 is 0.44.

2.5 Factor intensities and expenditure elasticities

We next study the correlation patterns between expenditure elasticities and factor intensities across sectors. Figure 3 plots the factor intensities in 1980 for low-, middle-, and high-skilled occupations in each sector against their corresponding expenditure elasticities.

Clear correlation patterns emerge. Expenditure-elastic sectors use high- and low-skilled occupations more intensively in production (correlations of 0.91 and 0.61, respectively). In contrast, there is no significant association (correlation of -0.08) between expenditure elasticity and the intensity of middle-skilled occupations.

These patterns are consistent with existing evidence linking expenditure elasticities to

[8]We describe the estimation procedure in detail in Section 4.2.

sectoral skill intensity. Prior work, using alternative estimation approaches, has largely relied on a binary distinction between high- and low-skill workers. By contrast, we consider three occupational groups and uncover a non-monotonic relationship across the skill distribution. Cravino and Sotelo (2019) and Andreolli et al. (2024) use consumer data to estimate income elasticities of demand for two broad sectoral aggregates. Cravino and Sotelo (2019) show that manufacturing is less intensive in high-skill workers than services. Andreolli et al. (2024) divide most of the economy into essential and non-essential sectors and show that the non-essential sector is relatively more intensive in hand-to-mouth agents. Leonardi (2015) shows that there is a positive association between the fraction of workers with some college employed in a sector and the education elasticity of the sector.[9] Caron et al. (2014) estimate sectoral income elasticities by combining cross-country trade data with input-output tables and assuming a CRIE demand system. They document a positive cross-sector association between income elasticities and the average (across 94 countries) share of skilled workers.

The heterogeneity of expenditure elasticities, together with the patterns of association between factor intensities and expenditure elasticities, opens the possibility for an income channel of labor-market polarization: as either average income or the dispersion of income across households 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 raises 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.

We next develop a framework that formalizes the income channel and use it to quantify the three channels to labor-market polarization.

3 Model

3.1 Firms

We consider an economy with $\mathcal{S} = \{1, \ldots, S\}$ sectors. Each sector $s \in \mathcal{S}$ produces output Y_{st} at time t by combining efficiency units of labor of three occupations $j \in \{H, M, L\}$, denoted by $\widetilde{X}_{jst}$, according to the Cobb-Douglas production function

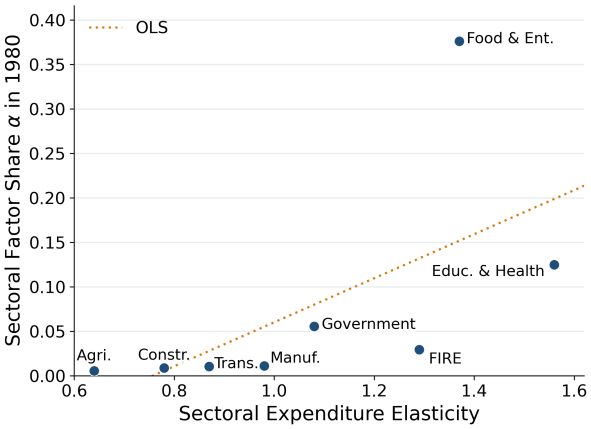
$$Y_{st} = \tilde{A}_{st} \prod_{j \in \{L,M,H\}} \widetilde{X}_{jst}^{\alpha_{jst}}, \tag{1}$$

where α_{jst} denotes the elasticity of value-added with respect to the efficiency units of labor of occupation j in sector s at time t. The three occupations correspond to High-, Middle- and Low-skill. We impose the constant-returns-to-scale constraint $\sum_{j \in \{L,M,H\}} \alpha_{jst} = 1$. $\tilde{A}_{st}$ is a

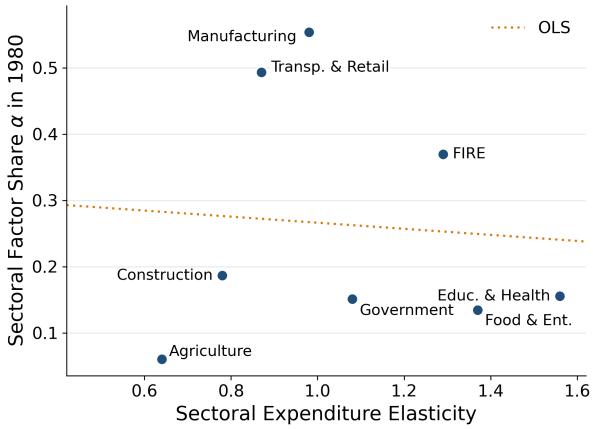
[9] He estimates the education elasticity of a sector by regressing household-level expenditure shares on a dummy for whether the head of the household has some college education.

Figure 3: Sectoral Factor Shares α and Expenditure Elasticities across Sectors

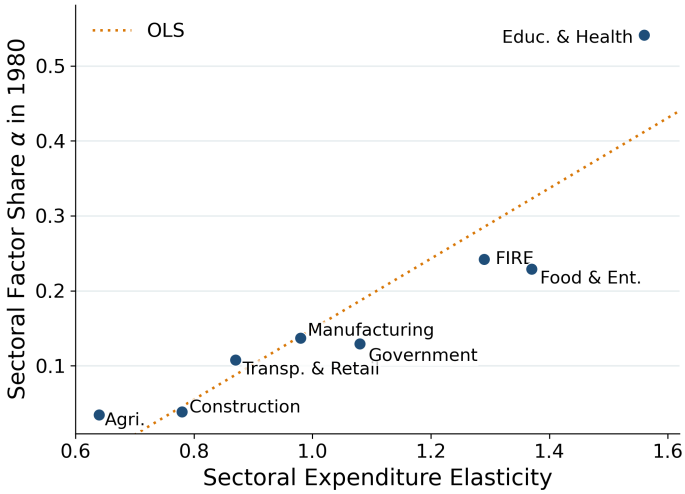
(a) Low-skill Occupations



(b) Middle-skill Occupations



(c) High-skill Occupations



Notes: The sectoral factor share α_{jst} of occupation j in sector s at time t is the wage bill of occupation j in sector s over the total wage bill in sector s. Employment is computed using hours worked from the Census/ACS, and wages are from the CPS (see Appendix A). Expenditure elasticities are estimated from household data using the Aguiar and Bils (2015) methodology (see Section 3.2).

sector-time productivity that captures both technological improvements and the contribution of capital that is not explicitly modeled in our baseline model. It is convenient to decompose $\tilde{A}_{st}$ into an aggregate component, A_t, and a sector-specific component, A_{st}, such that $\tilde{A}_{st} = A_t A_{st}$.

Firms are monopolistically competitive and charge a gross markup of ψ_t. The demand for efficiency units of labor of occupation j in sector s at time t is

$$\widetilde{w}_{jt}\widetilde{X}_{jst} = \frac{\alpha_{jst} P_{st} Y_{st}}{\psi_t}, \tag{2}$$

where $\widetilde{w}_{jt}$ is the price of an efficiency unit of labor of occupation j at time t, and P_{st} is the price of sectoral output,

$$P_{st} = \psi_t \frac{\prod_j \left(\frac{\widetilde{w}_{jt}}{\alpha_{jst}}\right)^{\alpha_{jst}}}{\tilde{A}_{st}}. \tag{3}$$

Note that $\widetilde{w}_{jt}$ differs from the average wage paid in occupation j, W_{jt}, as shown in Section 3.3.

3.2 Consumer preferences

There is a continuum of mass 1 of households indexed by $h \in \mathcal{H} = [0, 1]$. Household h has income E_{ht} and consumes a basket of sectoral goods and services, $\{c_{hst}\}_{s\in\mathcal{S}}$, whose value at prices $\{P_{st}\}_{s\in\mathcal{S}}$ satisfies $\sum_{s\in\mathcal{S}} P_{st}c_{hst} = E_{ht}$. Preferences over this basket are represented by a consumption index C_{ht}. The consumption index C_{ht} is implicitly defined by the non-homothetic CES aggregator

$$\sum_{s\in\mathcal{S}} \left(C_{ht}^{\varepsilon_s}\zeta_{st}\right)^{\frac{1}{\rho}} c_{hst}^{\frac{\rho-1}{\rho}} = 1, \tag{4}$$

where $\rho \in (0, 1)$ is the elasticity of substitution between consumption across the goods and services of different sectors, and the parameters $\varepsilon_s > 0$ govern the strength of non-homotheticity across sectors.[10] When ε_s is the same across sectors, preferences collapse to the standard homothetic CES.[11] As shown in Comin et al. (2021), we can normalize, without loss of generality, the taste parameter, $\zeta_s = 1$, and the income elasticity parameter, $\varepsilon_s = 1$, for one s. We choose government services to be the normalized sector.

Household h's optimal demand for good s is

$$c_{hst} = \zeta_{st}\left(\frac{P_{ht}}{P_{st}}\right)^{\rho} C_{ht}^{\varepsilon_s+\rho}, \tag{6}$$

[10] It is instructive to express this utility function as

$$C_{ht} = \left[\sum_s \left(\zeta_{st}C_{ht}^{\varepsilon_s+\rho-1}\right)^{\frac{1}{\rho}} c_{hst}^{\frac{\rho-1}{\rho}}\right]^{\frac{\rho}{\rho-1}} \tag{5}$$

where it is clear that this is a generalization of the homothetic CES aggregator, where the weight given to each sector varies with the consumption index C_{ht} and it is controlled by the parameter ε_s.

[11] In Appendix B, we generalize the preferences by using a nested non-homothetic CES specification. This formulation allows the elasticity of substitution to differ within and between broader sectoral aggregates (e.g., manufacturing, services, agriculture). We estimate the system and show that the within and between elasticities of substitution are quite similar implying that the standard non-homothetic CES is a good approximation to the data.

where $P_{ht} \equiv \frac{E_{ht}}{C_{ht}}$ is a household-specific price index given by

$$P_{ht} = \left[\sum_{s\in\mathcal{S}} \left(\zeta_{st} P_{st}^{1-\rho}\right)^{\theta_s} \left(x_{hst} E_{ht}^{1-\rho}\right)^{1-\theta_s}\right]^{\frac{1}{1-\rho}}, \tag{7}$$

with $x_{hst} \equiv P_{st} c_{hst}/E_{ht}$ denoting the expenditure share in sector s, and $\theta_s \equiv (1-\rho)/\varepsilon_s$. The associated expenditure function to the household problem is $E_{ht}^{1-\rho} = \sum_{s\in\mathcal{S}} \zeta_{st} C_{ht}^{\varepsilon_s} P_{st}^{1-\rho}$.

Aggregating (6) across households, we obtain the aggregate demand for sectoral output s,

$$C_{st} = \zeta_{st} \int_0^1 C_{ht}^{\rho+\varepsilon_s} \left(\frac{P_{ht}}{P_{st}}\right)^{\rho} dh. \tag{8}$$

Note that the curvature of aggregate demand varies across sectors. This observation is important because it opens the possibility for the dispersion of expenditures across households to affect the sectoral composition of the economy. In Appendix C, we use a second-order Taylor expansion of household sectoral demand $C_{hs}(E_h)$ around the average household expenditure $\bar{E}$ (time subscripts suppressed) to obtain that aggregate demand satisfies

$$\frac{C_s}{\bar{C}_s(\bar{E})} \approx 1 + \tfrac{1}{2}\left[\xi_s(\bar{E})\left(\xi_s(\bar{E}) - 1\right) + \bar{E}\,\xi_s'(\bar{E})\right] \frac{\mathrm{Var}(E_h)}{\bar{E}^2}, \tag{9}$$

where $\bar{C}_s(\bar{E})$ is the demand of a representative household with expenditure $\bar{E}$, $\mathrm{Var}(E_h)$ denotes the variance of expenditures across households, and $\xi_s(E)$ is the expenditure elasticity of demand for sector s evaluated at E. The bracketed term in Equation (9) governs the sign of the dispersion effect: when positive, expenditure dispersion raises aggregate demand for sector s above the representative-household benchmark $\bar{C}_s(\bar{E})$; when negative, the opposite occurs.

For non-homothetic CES preferences, Appendix C shows that this term is positive when the sector's expenditure elasticity ξ_s exceeds a threshold strictly above one, and negative for sectors with below-average expenditure elasticity. Thus, mean-preserving increases in the variance of expenditures across households raise aggregate consumption in sectors that are sufficiently income-elastic, and reduce it in sectors that are relatively income-inelastic.

3.3 Occupational choice

Each household is endowed with one unit of labor, which it supplies inelastically to one of the three occupations, high-, middle- and low-skill, $j \in \{H, M, L\}$. Household h draws a vector of occupation-specific efficiency units $\boldsymbol{\eta}_h = (\eta_{hH}, \eta_{hM}, \eta_{hL})$, where η_{hj} is the number of efficiency units the household would supply if it worked in occupation j. Given the vector of

wages per efficiency unit in the three occupations $\widetilde{\mathbf{w}}_t = (\tilde{w}_{Ht}, \tilde{w}_{Mt}, \tilde{w}_{Lt})$, household h chooses the occupation that maximizes labor income,

$$o_{ht} = \underset{j \in \{H,M,L\}}{\arg\max} \{\eta_{hj} \widetilde{w}_{jt}\}. \tag{10}$$

Household labor income is therefore $WB_{ht} = \max_{j \in \{H,M,L\}} \{\eta_{hj} \widetilde{w}_{jt}\}$.

Efficiency units are independently drawn across occupations from lognormal distributions.[12] We denote the mean and standard deviation of $\log \eta_j$ by μ_{jt} and σ_{jt}.[13] We interpret these idiosyncratic draws of efficiency units as a reduced-form mapping from a vector of abilities to worker productivity across occupations. The occupational-choice block should therefore be interpreted as a parsimonious reduced-form representation of occupational selection and effective labor supply, not as a structural model of education, training, or skill acquisition. In particular, changes in σ_{jt} stand in for any force that changes the effective supply of labor to occupation j, including changes in educational attainment, occupational training, occupational sorting, and changes in the task content of jobs. Our calibration does not separately identify these underlying forces, it uses the occupational-choice block to discipline the equilibrium supply response across occupations conditional on the observed changes in wages and hours-worked shares.

Define the set of draws that lead a household to select occupation j as

$$\mathcal{A}_{jt}(\widetilde{\mathbf{w}}_t) = \{\boldsymbol{\eta} : \eta_j \widetilde{w}_{jt} \geq \eta_k \widetilde{w}_{kt} \text{ for all } k \in \{H, M, L\}\}. \tag{11}$$

Let $F_t(\boldsymbol{\eta})$ denote the joint distribution of efficiency units across occupations. The share of households working in occupation j is then

$$\pi_{jt} = \int_{\mathcal{A}_{jt}(\widetilde{\mathbf{w}}_t)} dF_t(\boldsymbol{\eta}), \tag{12}$$

[12] With independent draws, the calibrated model matches the average wages of high- and low-skill workers relative to middle-skill workers, as well as the cross-sectional variance of log wage income in each period. This specification leaves within-occupation wage dispersion as an untargeted moment. As discussed below, the model produces dispersion levels in the same ballpark as the data, and matches the evolution of within-occupation wage dispersion reasonably well for high- and middle-skill workers, but not for low-skill workers. In Section 5.6, we develop an extension with correlated draws that can match these within-occupation dispersion, and show that our main results go through in this richer environment. A richer Roy-model structure with correlated multi-dimensional selection, explicit skill-acquisition costs, or idiosyncratic preferences for occupations would be valuable extensions; Erosa et al. (2025) make progress in this direction.

[13] Lognormal 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 distribution in the max-stable family as described in Lind and Ramondo, 2018) would imply the counterfactual prediction that average wage per worker is equalized across occupations. Papers using Fréchet distributions in similar settings therefore typically introduce unobserved costs or worker attributes; see, e.g., Galle et al. (2017).

and the supply of efficiency units to occupation j is

$$\widetilde{X}_{jt} = \int_{\mathcal{A}_{jt}(\widetilde{\mathbf{w}}_t)} \eta_j \, dF_t(\boldsymbol{\eta}). \tag{13}$$

The wage bill paid to workers in occupation j, WB_{jt} is

$$WB_{jt} = \widetilde{w}_{jt}\widetilde{X}_{jt}. \tag{14}$$

Hence, the average wage paid in occupation j is

$$W_{jt} = \frac{WB_{jt}}{\pi_{jt}} = \widetilde{w}_{jt}\frac{\widetilde{X}_{jt}}{\pi_{jt}}. \tag{15}$$

This expression makes clear why the wage per efficiency unit, $\widetilde{w}_{jt}$, differs from the average wage, W_{jt}: the latter also reflects the average efficiency units of the workers who select into occupation j, $\widetilde{X}_{jt}/\pi_{jt}$.

3.4 Household expenditures

Household expenditure E_{ht} equals total household income which is the sum of labor income WB_{ht} and profit income Z_{ht}:

$$E_{ht} = WB_{ht} + Z_{ht}. \tag{16}$$

Profits generated by firms are distributed unevenly across households. A fraction $(1-p)$ of households does not receive any profit income. The remaining share p receives profit income that is governed by a Pareto distribution that we calibrate to match the income share of the top-p percent of the distribution. Formally, let Z_{ht} denote the profit income accrued by household h at time t.

$$Z_{ht} = \begin{cases} 0 & \text{with probability } (1-p) \\ WB_{ht} \cdot D_{ht} & \text{with probability } p \end{cases} \tag{17}$$

where D_{ht} is a random variable that follows a Pareto distribution with lower bound 1 and shape parameter κ_t.[14]

[14]Note that the condition that the total profit income equals the aggregate profits in the economy imposes a linear relationship between the gross markup and the Pareto parameter:

$$\int_0^1 Z_{ht} dh = p \cdot WB_t \frac{\kappa_t}{\kappa_t - 1} = (\psi_t - 1) \cdot WB_t.$$

This implies that $\psi_t = 1 + \frac{p\kappa_t}{\kappa_t - 1}$. Throughout the paper, we assume that this condition holds. Alternatively,

3.5 Competitive equilibrium

The economy is static within each date: households do not save and firms make only contemporaneous production decisions. We can therefore solve the model period by period and define the competitive equilibrium for an arbitrary date t.

Given the distributions of efficiency unit draws, F_t, and profit income draws, D_{ht}, a competitive equilibrium at date t consists of sectoral prices $\{P_{st}\}_{s\in\mathcal{S}}$, wages per efficiency unit $\{\widetilde{w}_{jt}\}_{j\in\{H,M,L\}}$, household consumption choices $\{c_{hst}\}_{h\in\mathcal{H},s\in\mathcal{S}}$, occupational choices $\{o_{ht}\}_{h\in\mathcal{H}}$, and sectoral production and labor allocations $\{Y_{st}, \widetilde{X}_{jst}\}_{s\in\mathcal{S},j\in\{H,M,L\}}$ such that:

i. Household expenditures E_{ht} are given by (16), where profit income is given by (17).

ii. Given prices and final household expenditure E_{ht}, each household chooses consumption to maximize (4) subject to the budget constraint $\sum_{s\in\mathcal{S}} P_{st}c_{hst} = E_{ht}$.

iii. Given wages per efficiency unit, each household chooses its occupation according to (10).

iv. Firms minimize costs taking prices and wages as given. Their labor demand and sectoral prices satisfy (2) and (3), with corresponding output Y_{st} given by (1).

v. Goods markets clear:
$$\int_{\mathcal{H}} c_{hst}\, dh = Y_{st}, \qquad \forall s \in \mathcal{S}.$$

vi. Labor markets clear. For each occupation, the wage bill implied by firms' labor demand equals the wage bill generated by the Roy occupational supply:
$$\sum_{s\in\mathcal{S}} \frac{\alpha_{jst}P_{st}Y_{st}}{\psi_t} = \widetilde{w}_{jt}\widetilde{X}_{jt},$$
where $\widetilde{X}_{jt}$ is given by (13). Since firm demand satisfies $\widetilde{w}_{jt}\widetilde{X}_{jst} = \frac{\alpha_{jst}P_{st}Y_{st}}{\psi_t}$, this condition is equivalent to $\sum_{s\in\mathcal{S}} \widetilde{X}_{jst} = \widetilde{X}_{jt}$.

4 A Quantitative Model

Next, we calibrate the model to obtain a quantitative framework that allows us to study its implications. We adopt a hybrid approach: we estimate the preference parameters directly

we could scale the Pareto shock so that this condition holds automatically. This would break the dependence of ψ_t and κ_t providing an additional degree of freedom to match some other moment in the data. In the computation, we use a quantile function of a Pareto type-I distribution with lower bound 1 and tail parameter κ_t, $Q(u_m) = (1-u_m)^{-1/\kappa_t}$. We implement n_D equal-probability midpoint quantiles, $u_m = (m-1/2)/n_B$, with $m = 1, \ldots, n_D$.

using the Consumer Expenditure Survey (CEX), and calibrate the remaining objects to match key moments of the data in 1980 and 2016. The calibration separates the occupational-choice block from the profit income draws D_{ht}: the former is disciplined by labor-market moments, while the latter is disciplined by upper-tail household-income moments. In what follows, we describe the calibration strategy, explain how we estimate the preference parameters, and discuss the remaining parameter values.

4.1 Calibration Strategy

We calibrate the following objects $\{\{\zeta_{st}, \varepsilon_s, \alpha_{jst}, A_{st}\}_{s\in\mathcal{S}}, \rho, p, \kappa_t, A_t, \sigma_{jt}\}$ for $t \in \{1980, 2016\}$. We also solve for the vector of efficiency wages $\widetilde{\mathbf{w}}_t$. We normalize $\{\mu_{jt}\}$ for $j \in \{L, M, H\}$ in each period because efficiency units of labor are unit-free.[15] The calibration strategy is as follows:

i. The demand elasticities, ρ and $\{\varepsilon_s\}_{s\in\mathcal{S}}$ are estimated from the Consumer Expenditure Survey. We normalize $\varepsilon_s = 1$ for $s =$ government services.

ii. The sectoral occupation intensities, $\{\alpha_{jst}\}_{j\in\{L,M,H\},s\in\mathcal{S}}$, are set to match the wage-bill share of each of the three occupations in each of the eight sectors in 1980 and 2016.

iii. The occupational-choice block solves jointly for $\{\sigma_{jt}\}_{j\in\{L,M,H\}}$ and $\widetilde{\mathbf{w}}_t$ to match the aggregate occupation wage-bill targets constructed from point ii. above, the wage per worker of high- and low-skill occupations relative to middle-skill occupations, and the overall variance of log labor income in 1980 and 2016.

iv. We set $p = 0.05$ so that 5% of households receive profit income, and calibrate κ_t to match the income share of the top 5% of the distribution.[16]

v. The demand taste parameters, $\{\zeta_{st}\}_{s\in\mathcal{S}}$, are set to match sectoral share in value added in 1980 and in 2016. We normalize $\zeta_s = 1$ for $s =$ government services.

vi. Relative sectoral productivities, $\{A_{st}\}_{s\in\mathcal{S}}$, are set to match the relative sectoral prices in 1980 and 2016.

[15] To illustrate this claim, note that, given the wage bill for occupation j, WB_{jt}, the wage per efficiency unit is $\widetilde{w}_j = WB_j/\widetilde{X}_j$. If the occupational choices of agents do not change, increasing η_j by a factor $\kappa_j > 0$ would increase $\widetilde{X}_j$ by the same factor and reduce $\widetilde{w}_j$ by the same factor. As a result, household compensation and occupational choices would be unaffected. We also choose a numéraire for $\widetilde{\mathbf{w}}_t$ by normalizing the wage per efficiency unit in the middle-skill occupation, $\widetilde{w}_{Mt}$, to one in each period. Given this normalization, item iii. below matches the share of the total wage bill accruing to each occupation rather than its level, the aggregate level is then pinned down by A_t in item vii.

[16] Note that, a priori, there is no reason why the share of the population that receives profit income should coincide with the fraction of households used to define the top-income target.

vii. Aggregate productivity, A_t, is set to match the levels of expenditures per working-age person in 1980 and 2016.[17]

Discussion The calibration strategy has a useful block-recursive structure. The occupational-choice block is disciplined by production and labor-market moments only: sectoral value-added shares, sectoral occupation intensities, occupation wage bills, relative wages across occupations, and the log wage-bill dispersion target. Given these objects, the occupational-choice block determines occupational choices, efficiency-units supplied, occupation wage bills, average wages in each occupation, and the distribution of labor income WB_{ht} independently of the demand system.

The distribution of profit income is calibrated only after the occupational-choice block has been solved. Since Z_{ht} is occupation-orthogonal and independent of $\boldsymbol{\eta}_h$, it does not enter occupational sorting, occupation wage bills, or the labor-market moments used to identify $\{\sigma_{jt}\}_j$ and $\widetilde{\mathbf{w}}_t$ of the Roy model. Instead, it operates through the household expenditure distribution, which shapes sectoral demand under non-homothetic preferences. The demand block is then used to determine the taste parameters that rationalize sectoral value-added shares. The following lemma formalizes this separation; the proof is provided in Appendix D.

Lemma 1 (Occupational-choice block is an independent calibration block) *Conditional on the observed sectoral value-added shares, the sectoral occupation intensities* $\{\alpha_{jst}\}_{j\in\{L,M,H\},s\in\mathcal{S}}$*, and the labor-market moments used to discipline the occupational-choice block (Roy model), namely occupation wage bills, the wage per worker of high- and low-skill occupations relative to middle-skill occupations, and the variance of log labor income target, the occupational-choice block can be calibrated independently of the demand parameters* ρ*,* $\{\varepsilon_s\}_{s\in\mathcal{S}}$*,* $\{\zeta_{st}\}_{s\in\mathcal{S}}$*, and the distribution of profit income,* κ_t*.*

As we show below, Lemma 1 is relevant because it ensures that the values of the dispersion of productivity draws, aggregate TFP, relative sectoral TFP, and factor intensities required to match the targets in 1980 and 2016 are the same across different versions of the model featuring different demand systems.

Labor-only production function Our calibration matches aggregate expenditure, sectoral value added shares and the wage-bill of each occupation within each sector. However, it does not match the wage bills of each occupation observed in the data. This is the case because the model imposes a common labor share in value added across all sectors, which prevents us from matching sector-specific variation in wage bills. That is, in our model the

[17] A_t affects symmetrically sectoral prices, and it is set so that $E_t = \sum_s P_{st}C_{st}$ where C_{st} is given by (8).

Table 2: Occupation Wage-Bill Shares, Data vs. Constant Labor Share Construction

	1980		2016		2016–1980	
	Data	Const. Lab. sh.	Data	Const. Lab. sh.	Data	Const. Lab. sh.
High	30.5	28.8	46.5	49.9	17.8	19.4
Middle	62.5	64.5	45.3	41.4	-19.2	-21.1
Low	7.0	6.7	8.2	8.7	1.4	1.8

Notes: Entries are occupation wage-bill shares in percent. The last two columns report percentage-point changes from 1980 to 2016. Shares under the Constant Labor share columns are computed from Equation (18), while data shares are computed from Equation (19).

wage-bill share of occupation j is

$$\left(\frac{WB_{jt}}{WB_t}\right)_{Model} = \frac{\sum_s \alpha_{jst} VA_{st}}{\sum_j \sum_s \alpha_{jst} VA_{st}}, \tag{18}$$

while in the data it is

$$\left(\frac{WB_{jt}}{WB_t}\right)_{Data} = \frac{\sum_s \alpha_{jst}\beta_{st} VA_{st}}{\sum_j \sum_s \alpha_{jst}\beta_{st} VA_{st}}, \tag{19}$$

where β_{st} is the labor wage-bill share in value added for sector s. Because the capital share is positive and different across sectors, the distributions of value added and the wage bill across sectors will differ.

How significant is this deviation? Table 2 shows the wage-bill shares for each occupation in the data and constructed under the assumption of constant labor shares for 1980 and 2016, as well as the increments over the interval. Omitting the heterogeneity in β_{st} does not significantly alter the distributions of the wage bill across occupations in 1980, 2016 or their increments. Specifically, the increments in the wage-bill share for high- and low-skilled occupations in the model are 1.6 and 0.4 p.p. lower than in the data. These magnitudes are small relative to the actual increases in the wage-bill shares for these occupations (19.4 and 1.8 p.p., respectively). This rationalizes our modeling choice of abstracting from changes in the labor share to study labor-market polarization.[18,19]

[18]Since we match every period the wages of high- and low-skilled occupations relative to middle-skill occupations, and we slightly under-estimate the increase in their wage bills shares, we also slightly under-estimate the increase in the hours-worked shares of high- and low-skilled occupations. These are 11.6 vs. 13 p.p. for high and 2.8 vs. 3.4 p.p. for low-skilled occupations.

[19]It is nevertheless instructive to understand the sources of the gap between the changes in occupations' wage-bill shares in the model and the data. A priori, two channels could generate this discrepancy. First, initial occupation shares may be correlated with changes in the labor share across sectors. Second, initial labor shares may be correlated with changes in sectoral value-added shares. We find that the first channel is the main culprit. The correlation between the initial occupation share and the change in the labor share across sectors is 0.18 for high-skilled, -0.65 for middle-skill, and 0.84 for low-skilled occupations. By contrast,

4.2 Estimates of consumer preferences

We estimate the non-homothetic CES demand parameters: ρ, the elasticity of substitution across sectors, and $\{\varepsilon_s\}$, the sector-specific income-elasticity parameters. Together, they determine how changes in income and relative prices reallocate demand across sectors in the quantitative model. Since our estimation borrows extensively from previous literature (Aguiar and Bils, 2015 and Comin et al., 2021 in particular), we provide a relatively concise description of our data and estimation exercises in the main text and relegate the details to Appendix A.2.

Data Description We use household expenditure data from the Consumer Expenditure Survey (CEX) for 2000–2001. We follow Aguiar and Bils (2015) and Krueger and Perri (2006) in their data sample restrictions (e.g., studying only urban households). We convert the roughly 900 CEX expenditure categories into value-added consumption in eight broad sectors by mapping the detailed CEX expenditure categories to PCE lines and then to BEA's industries in the input-output tables, following the procedure in Buera et al. (2021) and Comin et al. (2021).[20] One limitation of the CEX data is that it only measures out-of-pocket expenditures. To mitigate the omission of partially publicly provided goods, we impute public expenditures (K-12 education, Medicare, and Medicaid) to the relevant CEX expenditure categories using geographic variation in per-pupil spending and hospital-referral-region health expenditures, as well as fund-management expenses that are subtracted from fund payouts and thus underreported in the CEX. Appendix A.2 provides full details on data construction and imputation procedures. Following Buera et al. (2021) and Comin et al. (2021), we supplement the household data with sectoral urban price series from the BLS and construct value-added price indices using the same input-output cost shares used for expenditures.[21] Before estimating the structural demand parameters, we also document in Appendix A.2 substantial cross-sector heterogeneity in expenditure elasticities using the Aguiar and Bils (2015) reduced-form approach.

the correlation between changes in value added shares and initial labor shares is only 0.05.

[20] The eight-sector aggregation balances economic content and modeling parsimony. It preserves the main heterogeneity in expenditure elasticities and occupation intensities while keeping the demand system simple.

[21] Our analysis follows the consumption value-added construction in Herrendorf et al. (2013), in which final-expenditure vectors are mapped into value-added expenditures using the total requirements matrix. Under the assumption of a Cobb-Douglas aggregator for this constructed composite, the normalized IO expenditure shares correspond to cost shares. The dual unit-cost function then implies that log changes in the corresponding value-added price index are cost-share-weighted averages of log changes in the corresponding final-expenditure price indices. This is a Stone price index with the IO expenditure shares as weights. The full construction, including a comparison with Leontief aggregation is in Appendix A.2.1.

Estimation Following Comin et al. (2021), the log-linear structure of the non-homothetic CES demand system allows us to invert the demand for a reference sector $\hat{s}$ to express the real consumption index of household h, C_{ht}, as a function of observables. Substituting this into the demand equations for all other sectors and parameterizing taste shifters with household demographics, we obtain the following estimating equation for all $s \neq \hat{s}$

$$\ln x_{hst} = \Gamma_s Z_h + \delta_{sr} + (1-\rho)\ln\left(\frac{P_{hst}}{P_{h\hat{s}t}}\right) + (1-\rho)(\varepsilon_s - 1)\ln\left(\frac{E_{ht}}{P_{h\hat{s}t}}\right) + \varepsilon_s \ln x_{h\hat{s}t} + u_{hst}. \quad (20)$$

where x_{hst} is household h's expenditure share in sector s, E_{ht} is total expenditure, p_{hst} and $p_{h\hat{s}t}$ are observed sector and reference-sector household-specific prices, $x_{h\hat{s}t}$ is reference-sector expenditure share, Z_h is a vector of demographic control dummies (age bins, number of earners, family size),[22] and δ_{sr} denotes region-by-sector fixed effects. Note that the system defines expenditure shares for $s \neq \hat{s}$ exclusively in terms of observables. We use this resulting system of $S-1$ equations in our estimation.[23] We also include the reduced-form expenditure elasticities from Aguiar and Bils (2015) as an additional set of moments.[24]

We estimate the system via GMM, following the instrumentation strategy in Comin et al. (2021) to deal with potential measurement error and endogeneity concerns. Household income levels and income quintile dummies serve as instruments for total expenditure, capturing permanent income while mitigating transitory measurement error. We instrument household relative prices with a Hausman-style relative-price instrument, which is constructed in two steps. First, for each sub-component of a sector, we compute the average price across regions

[22] Leonardi (2015) argues that the non-homotheticity in demand is associated to the level of education rather than to income. In his model, college-educated consumers have different preferences (i.e. taste parameters) than non-college-educated consumers. We have explored whether non-homotheticities are associated to income or education by including a dummy when the head of the household has college education in the vector Z_h. Note that, because Γ_s is sector-specific, this captures Leonardi's idea. Table A1 reports the estimates and shows that the introduction of education-specific taste parameters does not alter the expenditure-elasticities across sectors. In addition, we have re-estimated the reduced-form Aguiar-Bils regressions adding the same college dummy as an additional control, and we have also estimated the expenditure elasticities separately for samples of households with and without a college-educated head. The expenditure elasticities from these three exercises correlate with our baseline estimates at 0.99, 0.94, and 0.99, respectively. This suggests that the cross-sector dispersion in expenditure elasticities is not driven by college-educated households, but rather reflects income effects in demand, as specified in the non-homothetic CES formulation.

[23] Since any sector can be used as a reference sector in (20), we use all $S(S-1)$ moments obtained by taking each sector $s \in \mathcal{S}$ as a reference.

[24] We have that

$$\hat{\eta}_s = \sigma + (1-\sigma)\frac{\varepsilon_s}{\sum_{s' \in \mathcal{S}} \bar{x}^n_{s'} \varepsilon_{s'}} + v_s, \quad (21)$$

where $\hat{\eta}_s$ is the expenditure elasticity estimated using the Aguiar and Bils (2015) method, $\bar{x}^n_{st}$ is the sample average expenditure share in sector s, and v_s denotes an error term. These moments are not needed for identification, since equation (20) identifies ρ, $\{\varepsilon_s\}$. Rather, they discipline the nonlinear mapping from structural parameters to the implied expenditure elasticities $\{\eta_s\}$, which are central to the income channel, and keep the structural estimates tightly connected to the reduced-form evidence in Section 2.

Table 3: Estimated Expenditure and Demand Elasticities

	Nonhomothetic CES[2]			Reduced-Form[3]
	ρ	ε_s	ξ_s	ξ_s
Sectors	(1)	(2)	(3)	(4)
Education and Health Care		1.77	1.54	1.56
		(0.02)	(0.01)	(0.06)
Arts, Entertainment, Recreation and Food Services		1.37	1.30	1.31
		(0.03)	(0.01)	(0.03)
Agriculture		0.32	0.63	0.64
		(0.02)	(0.01)	(0.02)
Construction		0.48	0.73	0.78
		(0.01)	(0.01)	(0.02)
Retail, Wholesale Trade and Transportation		0.59	0.79	0.87
		(0.01)	(0.01)	(0.01)
Manufacturing		0.74	0.89	0.98
		(0.03)	(0.01)	(0.01)
Finance, Professional, Information, other services (excl. gov't)		1.28	1.29	1.29
		(0.01)	(0.01)	(0.01)
Government[1]		1.00	1.06	1.08
		(-)	(0.01)	(0.01)
Elasticity of Substitution	0.42			
	(0.02)			

Notes: Standard errors clustered at the household level shown in parentheses. Total number of households is 7,800. 1: Gov't sector is normalized to 1 in the demand estimation. 2: ξ_s in column (3) is the model-implied expenditure elasticity for the sample average household expenditures at the estimated parameters $\{\epsilon_s, \xi_s\}$, according to the formula $\hat{\xi}_s = \rho + (1-\rho)\varepsilon_s / \sum_{s' \in \mathcal{S}} \bar{x}^n_{s'} \varepsilon_{s'}$, where $\bar{x}^n_{s'}$ denotes the household average expenditure share in the sample. 3: ξ_s in column (4) corresponds to the expenditure elasticity estimated using the Aguiar and Bils (2015) specification.

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.

Results The first and second columns of Table 3 report the parameter estimates. The elasticity of substitution is $\rho = 0.42$ implying that sectors are complements. This coincides with the estimate in Goos et al. (2014), obtained using European data for 10 sectors from 1993-2010. The non-homotheticity parameters ε_s vary substantially across sectors: Education and Health Care has the highest value (1.77), while Agriculture has the lowest (0.32). These translate into expenditure elasticities, evaluated at sample-average expenditure shares and reported in column (3), ranging from 1.54 for Education and Health Care to 0.63 for Agriculture. The ordering is intuitive: certain services such as education, health, finance, professional services and information, arts and entertainment are expenditure-elastic, while goods-producing

sectors such as agriculture, construction, and manufacturing are expenditure-inelastic.As a cross-check, Column (4) reports reduced-form expenditure elasticities estimated using the specification in Aguiar and Bils (2015). These elasticities closely match the model-implied values in Column (3), providing independent validation of the structural estimates. Appendix A.2 presents robustness checks showing that the estimates are stable across alternative imputation assumptions and sample periods.

4.3 Discussion of the calibration parameter values

Before using the quantitative model, we discuss the parameter values obtained through the calibration. Table 4 reports the evolution of (i) the standard deviation of the productivity draw distributions for each occupation, (ii) the Pareto distribution parameter for profit income, (iii) the growth in relative sectoral TFP, and (iv) the growth in the taste parameters.[25]

Demand for efficiency units of each occupation is determined by the sectoral composition of value added and by the factor intensities in each sector (α_{jst}). The supply of efficiency units for each occupation is governed by the dispersion of the productivity draws across occupations, which determines both the fraction of workers selecting into each occupation and the average efficiency units they supply conditional on their occupation choice. Therefore, conditional on the demand for sectoral value added, the dispersion of productivity draws pins down relative wages across occupations.

Panel A in Table 4 shows that the dispersions of productivity draws are relatively stable between 1980 and 2016. The standard deviation of high-skill productivity draws is essentially unchanged, the standard deviation of middle-skill draws rises from 1.03 to 1.10, and the standard deviation of low-skill draws falls slightly from 1.38 to 1.34. Therefore, given demand and factor intensities, the model does not require large changes in the variance of productivity draws to match the evolution of relative wages and the variance of log labor income across households.

Panel A also reports the calibrated Pareto tail parameter that governs profit income draws. In 1980, the distribution of labor income implies a top 5% income share slightly above the target, so the Pareto tail is inactive (i.e., $\kappa_t = \infty$). In 2016, the estimated tail parameter is $\kappa_{2016} = 1.139$, implying a substantially fatter upper tail for profit income.[26,27]

[25] Aggregate TFP growth is computed as the weighted average of sectoral TFP growth, using 1980 value-added shares as weights.

[26] The gross markup ψ_t is not a free parameter. It is determined by the Pareto tail κ_t through the relationship $\psi_t = 1 + p\kappa_t/(\kappa_t - 1)$ derived in footnote 14. With $p = 0.05$, the calibrated tail generates model-implied gross markups of $\psi_{1980} = 1.05$ and $\psi_{2016} = 1.41$, in line with recent estimates of rising U.S. firm markups over this period.

[27] Appendix Table A3 reports that the model perfectly matches the variance of the log labor income in the data. Appendix A describes how we have constructed the target moments that reflect the variance of the log

Table 4: Calibrated Parameter Values

Panel A. Occupational-choice and Profit-Income Pareto Tail	1980	2016
High-skill productivity dispersion, σ_{Ht}	1.48	1.48
Middle-skill productivity dispersion, σ_{Mt}	1.03	1.10
Low-skill productivity dispersion, σ_{Lt}	1.38	1.34
Pareto tail profit income draw, κ_t	∞	1.139

Panel B. Sector-Level Parameters	Rel. TFP Growth Net of Agg. TFP	Taste-Parameter Change $\zeta_{s2016}/\zeta_{s1980}$
Education and Health	-0.49	1.22
Arts, Entertainment, Recreation, and Food	-0.19	1.62
Agriculture	1.79	1.35
Construction	-0.37	1.47
Retail, Wholesale, and Transportation	0.25	1.78
Manufacturing	0.46	1.22
Finance, Professional, Information, and Other Services	-0.28	2.14
Government	-0.44	1.00

Notes: Panel A reports the calibrated standard deviations of occupation-specific productivity draws for the occupational-choice block, and the calibrated Pareto tail parameter of the profit-income draws. The value $\kappa_{1980} = \infty$ denotes an inactive tail. Panel B reports sectoral TFP growth net of aggregate TFP growth and changes in sectoral taste parameters. Aggregate TFP growth is $A_{2016}/A_{1980} = 2.042$.

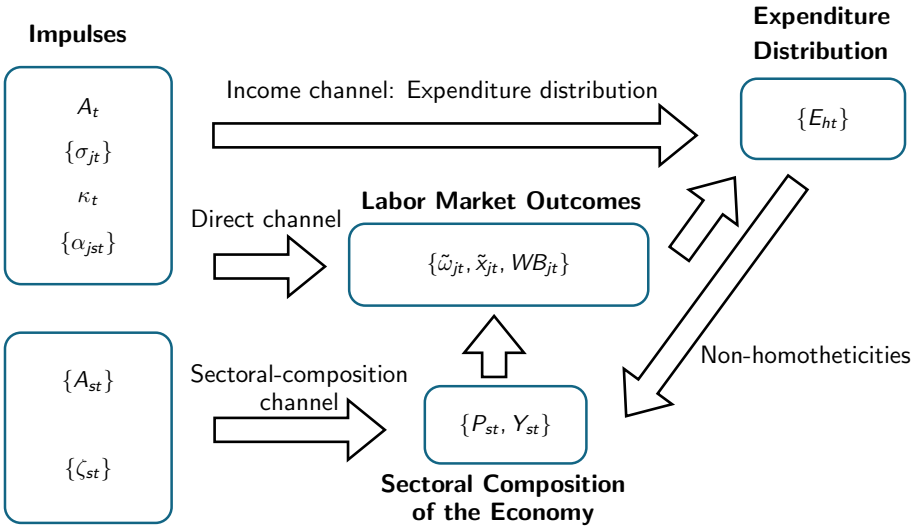
In our model, productivity growth arises from increases in aggregate productivity and from changes in the sectoral total factor productivity parameters, A_{st}. We identify sectoral TFP from the inverse of the residual in the BEA sectoral value-added deflators, given the equilibrium wages per efficiency unit of labor in each occupation. Panel B of Table 4 reports TFP growth in each sector from 1980 to 2016, expressed relative to the economy-wide average. We find that TFP growth was above average in agriculture, manufacturing, and retail, wholesale and transportation, and below average in the remaining sectors.

To exactly match the sectoral composition of value added in both 1980 and 2016, we allow the taste parameters, ζ_{st}, to vary over time. The second column in Panel B of Table 4 shows that the sectors requiring the largest increases in the taste parameter to match the observed changes in value-added shares are finance, professional, information, and other services; retail, wholesale, and transportation; and arts, entertainment, recreation, and food services. Government services is the normalized sector, so its taste parameter equals one by construction.

Finally, we also study the evolution of the within-occupation dispersion of log labor income, a moment the model does not target. Nevertheless, the model does a reasonable job matching

labor income, the top income share.

Figure 4: Drivers and channels of labor market polarization



both the levels and the growth of the within-occupation dispersion in log labor income (see Table A2 in the Appendix).[28] The model-implied 2016/1980 ratios are 1.051, 1.018, and 1.015, compared with 1.092, 0.999, and 0.946 in the data.[29]

5 Analysis

In this section, we put the quantitative framework to work. First, we develop a conceptual framework that highlights the channels through which the different impulses affect labor-market outcomes. Second, to operationalize this framework, we present two alternative versions of the model that progressively shut down the income channel. Third, we use counterfactuals to assess the impact of each individual shock on labor-market outcomes. Fourth, we examine the effects of combinations of shocks across the different versions of the model. Fifth, we use these exercises to quantify the three channels of labor-market polarization: the direct, sectoral-composition, and income channels. Sixth, we assess the robustness of the income channel to alternative model specifications, calibrations, and quantification approaches. Seventh, we discuss policies that influence the income and sectoral-composition channels.

5.1 Conceptual framework

Figure 4 presents a conceptual framework that illustrates how the different impulses affect labor-market outcomes; throughout the analysis, an impulse denotes a calibrated change in one or more economic primitives between 1980 and 2016. In our model, six impulses shape labor markets. They correspond to changes in: (1) the occupation intensities α_{jst}, (2) the dispersion of productivity draws σ_{jt}, (3) aggregate TFP A_t, (4) the Pareto tail for profit income draws κ_t, (5) relative sectoral TFP A_{st}, and (6) taste parameters ζ_{st}.

Changes in factor intensities directly affect labor-market outcomes by shifting the relative demand for different occupations. Similarly, changes in the dispersion of productivity draws and in aggregate TFP directly influence both the supply of efficiency units of labor across occupations and their demand. These direct effects of impulses on labor-market outcomes are represented in Figure 4 by the arrow labeled "Direct channel."

Changes in relative sectoral TFP and in the taste parameters do not directly affect relative labor demand or supply. However, as shown by Lee and Shin (2017) and Bárány and Siegel (2018) they influence labor-market outcomes indirectly by altering the sectoral composition of the economy and, because occupation intensities differ across sectors, the demand for labor across occupations. This channel is represented in Figure 4 by the arrow labeled "Sectoral composition channel."

Changes in occupation intensity, the variance of productivity draws, aggregate TFP, and the Pareto tail for profit income draws also affect labor market outcomes indirectly by changing the expenditure distribution across households. With non-homothetic preferences, increases in the average expenditure or mean-preserving spreads in the distribution of expenditure raise the value-added share of income elastic sectors. These sectors are more intensive in high- and low-skilled labor, so these impulses indirectly increase the relative demand for these occupations. This channel is represented in Figure 4 by the arrow labeled "Income channel: Expenditure distribution."

As this framework illustrates, studying the contributions of each impulse to labor-market outcomes in the baseline model does not reveal the magnitude of a channel, because multiple impulses contribute to each channel and a given impulse may operate through several channels. Instead, we quantify the three channels by first developing a version of the model with homothetic preferences, in which the income channel is shut down because changes in the expenditure distribution do not affect sectoral composition or labor-market outcomes. With the baseline and homothetic versions in hand, we quantify the channels as follows:

[28] In 1980, the model-implied standard deviations are 1.002, 0.910, and 0.897 for high-, middle-, and low-skill occupations, respectively, compared with 0.841, 0.897, and 0.935 in the data.

[29] In Section 5.6, we explore the robustness of our findings to allowing productivity draws across occupations to be correlated to match the observed within-occupation dispersion in wages.

- **Direct channel:** the effect of applying the shocks that directly influence labor supply and demand and/or the distribution of expenditures across households—changes in factor intensities, the variance of productivity draws, the Pareto tail parameter for profit income, and aggregate TFP— in the homothetic model.

- **Sectoral-composition channel:** the effect of shocks that modify the sectoral composition of the economy but do not directly affect the income distribution—relative sectoral TFP and taste parameters—in the baseline model.

- **Income channel:** the differential effect of shocks that directly influence labor supply and demand and/or the distribution of expenditures across households—changes in factor intensities, the variance of productivity draws, the Pareto tail parameter for profit income, and aggregate TFP—in the baseline model relative to the homothetic model.

A few points are worth stressing. First, because the direct effect is measured in the homothetic model, shocks that only affect the distribution of expenditures (e.g., κ_t) have no impact on this channel. Second, for shocks that directly affect labor-market outcomes (i.e., α_{jst}, σ_{jt}, and A_t) only their direct impact is reflected, because their impact on the expenditure distribution has no consequences in the homothetic model. Third, because the shocks included in the direct and income channel are the same, and the mechanisms underlying the direct channel are identical in the baseline and homothetic model, measuring the differential effect of these shocks in the two models isolates the effect of these shocks operating through changes in the expenditure distribution. Fourth, it is important that the sectoral-composition channel is evaluated in the baseline model because, as implied from Lemma 1, the taste parameters differ across model versions. Finally, in Section 5.6 we develop an alternative accounting approach that allows the income channel to reflect the (minor) changes in the expenditure distribution induced by changes in relative sectoral TFP.

5.2 Model variants without the income channel

We next develop two variations of the baseline model that progressively shut down the income channel. The first version removes heterogeneity in consumption choices across households by assuming a representative consumer with non-homothetic preferences. The second version removes non-homotheticities in demand by replacing the non-homothetic aggregator with a homothetic one. We just formalize the aspects that differ from the baseline.

Representative consumer with non-homothetic preferences

In the representative consumer economy, consumption choices are made by a representative consumer that collects all the income in the economy and that has preferences defined by

$$\sum_{s\in\mathcal{S}} \left(\left(C_t^R\right)^{\varepsilon_s} \zeta_{st}^R\right)^{\frac{1}{\rho}} C_{st}^{R\,\frac{\rho-1}{\rho}} = 1, \tag{22}$$

where C_t^R is the consumption index of the representative consumer, C_{st}^R is his consumption of the good or service from sector s, and ζ_{st}^R is the taste parameter. Note that we allow the taste parameters to differ from the baseline to match the sectoral composition of value added in the data (as discussed below).

The resulting consumption choices are:

$$C_{st}^R = \zeta_{st}^R \left(\frac{P_t^R}{P_{st}}\right)^{\rho} (C_t^R)^{\varepsilon_s+\rho}, \tag{23}$$

where $C_t^R = \frac{E_t}{P_t^R}$, E_t is aggregate expenditure, and P_t^R is the price deflator for the representative consumer, given by

$$P_t^R = \left[\sum_{s\in\mathcal{S}} \left(\zeta_{st}^R P_{st}^{1-\rho}\right)^{\theta_s} \left(x_{st}^R E_t^{1-\rho}\right)^{1-\theta_s}\right]^{\frac{1}{1-\rho}}, \tag{24}$$

with $x_{st}^R \equiv P_{st} C_{st}^R / E_t$.

Homothetic preferences

In the homothetic version of the model, preferences are defined by the standard CES aggregator (Dixit and Stiglitz, 1977). In particular, the consumption index of household h is

$$C_{ht}^H = \left(\sum_{s\in\mathcal{S}} \left(\zeta_{st}^H\right)^{\frac{1}{\rho}} \left(c_{hst}^H\right)^{\frac{\rho-1}{\rho}}\right)^{\frac{\rho}{\rho-1}}. \tag{25}$$

Aggregate consumption in sector s is

$$C_{st}^H = \zeta_{st}^H \left(\frac{P_t^H}{P_{st}}\right)^{\rho} C_t^H, \tag{26}$$

where $C_t^H = \int_0^1 C_{ht}^H dh$ is the aggregate consumption index, and $P_t^H = \left[\sum_s \zeta_{st}^H P_{st}^{1-\rho}\right]^{\frac{1}{1-\rho}}$ is the price deflator. Note that this deflator is the same for all households, and that, unlike the non-homothetic baseline, the homothetic economy admits a representative consumer.

Recalibration of alternative versions

Lemma 1 implies that, given $(\{\varepsilon_s\}, \rho)$, the values of $(\{\sigma_{jt}\}, \{\alpha_{jst}\}, \{A_{st}\}, \kappa_t, A_t)$ required to match relative wages, hours-worked shares, wage-bill shares across occupations and sectors, aggregate expenditure, the variance of log labor income across households, the top-income share and sectoral prices are identical in the two alternative formulations and in the baseline model. By contrast, we recalibrate the taste parameters $\{\zeta_{st}^R\}$ and $\{\zeta_{st}^H\}$ to match the 1980 and 2016 value-added shares across sectors in the representative-consumer and homothetic versions of the model.

5.3 Contribution of impulses to labor-market polarization in the baseline model

Next, we use model counterfactuals to quantify how each impulse contributes to the evolution of labor market outcomes. Table 5 reports the labor market outcomes in 2016 from four counterfactual exercises, each shocking our 1980 calibrated economy with a different impulse or group of impulses one at a time. For comparison, it also reports the outcomes produced when all shocks are applied in 2016. To shed light on the mechanism, the Table also reports the fraction of the observed change in the sectoral value-added shares accounted for by each impulse, and the covariance between the change in the sectoral distribution of value added induced by the impulse and the initial intensity of high- and low-skilled workers.

Changing α_{jst} Changing the sectoral factor intensities to their 2016 levels directly shifts the relative labor demand of high- and low-skilled occupations. This raises both the relative wages and the share of hours worked for workers in these occupations.

In particular, changing α_{jst} fully accounts for the increase in the relative wage of high-skilled occupations (20.5 vs. 19.5 log points) and low-skilled occupations (10.75 vs. 3.9 log points). Changes in factor intensity also explain 70.4% of the increase in the high-skilled hours-worked share (8.1 p.p. vs 11.6 p.p.) and 40% of the increase for low-skilled workers (1.1 p.p. vs. 2.8 p.p.). As a result, the change in factor intensities accounts for 78% of the increase in the wage-bill share of high-skilled workers (13.8 p.p. vs. 17.7 p.p.) and 56% of the increase for low-skilled workers (0.8 p.p. vs. 1.4 p.p.).

Changing σ_{jt}, κ_t and A_t This impulse has no effect on the relative wage of high-skilled workers and reduces by 3.7 log points the relative wage of low-skilled workers. It does, however, account for substantial shares of the observed occupational reallocation between 1980 and 2016. Specifically, it explains 31% and 24% of the increase in the high-skilled hours-

Table 5: Labor-Market Polarization in Model Counterfactuals: Individual Changes in Fundamentals

	All Impulses			Counterfactual Changes Relative to 1980				
	1980	2016	Change	Factor Intensities α_{jst}	Roy Dispersion Pareto Tail and Agg. TFP $\sigma_{jt} + \kappa_t + A_t$	Relative Sectoral TFP A_{st}	Taste Parameters ζ_{st}	Sectoral Composition $A_{st} + \zeta_{st}$
Panel A. Relative Wages			$100 \times \Delta \log$ Relative Wages, *2016-1980*					
High/Middle	1.26	1.53	19.4	20.5	0.0	2.8	-4.4	-2.1
Low/Middle	0.77	0.80	3.9	10.7	-3.8	3.9	-6.8	-3.0
Panel B. Employment Shares (%)			*Percentage-point differences, 2016-1980*					
High	23.7	35.3	11.6	8.2	3.6	0.9	-1.4	-0.7
Middle	67.2	52.8	-14.4	-9.3	-6.1	-1.6	2.6	1.2
Low	9.1	11.9	2.8	1.1	2.5	0.8	-1.3	-0.6
Panel C. Wage-Bill Shares (%)			*Percentage-point differences, 2016-1980*					
High	28.7	46.5	17.7	13.8	4.3	1.5	-2.4	-1.2
Middle	64.5	45.4	-19.2	-14.6	-5.9	-2.3	3.6	1.7
Low	6.7	8.2	1.4	0.8	1.6	0.8	-1.2	-0.6
Panel D. Sectoral-Composition Metrics			*Sectoral-composition statistics*					
$100 \times \mathrm{Cov}(\Delta sh_s^{\mathrm{all}}, \Delta sh_s^{\mathrm{cf}}) / \mathrm{Var}(\Delta sh_s^{\mathrm{all}})$			100	-1	44	44	19	62
$100 \times \mathrm{Cov}(\Delta sh_s, \alpha_H)$			0.44	0.05	0.62	0.24	-0.35	-0.18
$100 \times \mathrm{Cov}(\Delta sh_s, \alpha_L)$			0.16	0.05	0.22	0.12	-0.18	-0.09

Notes: Under "All Impulses," the first two columns report the 1980 and 2016 calibrated equilibria. The third column reports the 1980–2016 change: Panel A reports $100 \times \Delta \log(W_i/W_M)$ for $i = H, L$, while Panels B and C report percentage-point changes in shares. Each counterfactual column reports the corresponding change relative to the 1980 baseline. The first row of Panel D reports the projection statistic $100 \times \mathrm{Cov}(\Delta sh_s^{\mathrm{all}}, \Delta sh_s^{\mathrm{cf}}) / \mathrm{Var}(\Delta sh_s^{\mathrm{all}})$, where $\Delta sh_s^{\mathrm{all}}$ is the all-impulses change in sectoral value-added shares and $\Delta sh_s^{\mathrm{cf}}$ is the change induced by the counterfactual. The remaining Panel D rows report the covariance of the induced sectoral reallocation with 1980 high- and low-skill factor intensities.

worked and wage-bill shares from 1980 to 2016; 42.6% and 30.7% of the reduction in the hours-worked and wage-bill shares of middle-skill occupations; and 90.7% and 107% of the increase in the hours-worked and wage-bill shares of low-skill occupations.

Why does this impulse have an asymmetric effect on the evolution of relative wages and hours-worked shares? The changes in σ_{jt}, κ_t and A_t operate through two channels. First, σ_{jt} directly raises the relative supply of high- and low-skilled workers, putting downward pressure on the relative wages of these occupations while raising their hours-worked shares.

Second, this impulse fully accounts for the increase in average real expenditure and has a sizable effect on its dispersion.[30] The rise in both the mean and the dispersion of real expenditure shifts value added towards income-elastic sectors. Because income-elastic sectors are more intensive in high- and low-skill labor, this indirect force raises the relative demand for high- and low-skilled workers, mitigating the negative effect of the direct increase in relative supply on relative wages.

The reallocation effect of this impulse is substantial. It accounts for 44% of the observed change in the sectoral composition of value-added from 1980 to 2016, and the covariance between the change in the sectoral value added share and the 1980 intensities of high- and low-skilled occupations is 0.62/100 and 0.22/100, respectively.

Changing A_{st} This impulse increases the relative wages and the hours-worked shares of high- and low-skilled workers. It raises the relative wages of high- and low-skilled workers by 2.8 and 4 log points, respectively, and increases their hours-worked share by 0.9 and 0.8 p.p., respectively.

It is remarkable that this impulse has no direct effect on labor markets. Therefore, all of its effects are indirect—operating, as we show in the next section, through changes in the sectoral composition of the economy. The sectoral reallocation induced by this impulse accounts for 44% of the variation in the observed changes in the sectoral value-added shares. These changes in sectoral composition covary strongly with the 1980 intensity in high- and low-skilled workers (0.24/100 and 0.12/100, respectively). As a result, the impulse indirectly raises the relative demand for high- and low-skilled workers, leading to increases in both their relative wages and their hours-worked shares.

Changing ζ_{st} Changes in taste parameters do not contribute to explaining labor-market polarization. They have no direct effect on labor markets and affect labor market outcomes only indirectly by altering the sectoral composition of the economy. Unlike shocks to sectoral

[30]Real expenditure is nominal expenditure divided by the price index that results from weighting sectoral prices by value-added shares. Real expenditure in 2016 is 1.03 in the model vs. 1 in the data.

TFP, this impulse induces a relative expansion of sectors that are intensive in middle-skill occupations, thereby reducing the relative demand for high- and low-skilled workers.

5.4 Comparison across model versions

We next compare how combinations of impulses affect labor-market outcomes in the baseline economy and in the two alternative model economies. We begin with a combination of the impulses that directly influence labor markets and/or the distribution of income across households: the variance of productivity draws (σ_{jt}), factor intensities (α_{jst}), the Pareto tail parameter (κ_t), and aggregate TFP (A_t). We then examine some of these shocks individually and conclude with the impulse to relative sectoral TFP (A_{st}). The results are reported in Table 6.

Labor market and income distribution combination We first compare the effects of simultaneously changing σ_{jt}, κ_t, α_{jst} and A_t in the baseline and homothetic models. The impact of this combination of impulses on labor-market outcomes is considerably smaller in the homothetic economy than in the baseline. In the baseline model, the relative wages of high- and low-skilled occupations increase by 21.2 and 6.7 log-points, respectively, whereas in the homothetic version, the relative wage of high-skilled workers increases by 13 log points and that of low-skilled workers decreases by 1.4 log points. Similarly, the hours-worked share of high- and low-skilled workers increase by 12.1 p.p. and 3.4 p.p., respectively, in the baseline model, compared with 9.1 p.p. and 2 p.p. in the homothetic version.

Why do the same impulses generate different effects in the baseline and homothetic versions of the model?

This combination of impulses affects labor-market outcomes through three channels. The first reflects the direct effects on labor supply and labor demand: the variance of productivity draws drives the relative supply of high- and low-skilled occupations, while factor intensities affect their relative demand. These direct effects, however, are identical across all versions of the model

The second channel operates through changes in the sectoral composition of the economy induced by changes in relative prices and taste parameters. However, this mechanism is small, as reflected in the small covariances between changes in sectoral value-added shares and high- and low-skilled intensities across sectors in the homothetic model (0.02/100 and 0.014/100, respectively).

The third channel is the income channel. This combination of impulses increases both the average real consumption index in the economy and the dispersion of real consumption indices

Table 6: The Income Channel to Labor-Market Polarization, 1980–2016

	All Impulses	Counterfactuals: Mechanism(s) Active											
		Joint Labor-Expenditure Impulses $\sigma_{jt} + \kappa_t + A_t + \alpha_{jst}$			Roy Dispersion Pareto Tail and Agg. TFP $\sigma_{jt} + \kappa_t + A_t$			Factor Intensities α_{jst}			Relative Sectoral TFP A_{st}		
		Base.	Rep. Cons.	Homo-thetic	Base.	Rep. Cons.	Homo-thetic	Base.	Rep. Cons.	Homo-thetic	Base.	Rep. Cons.	Homo-thetic
Panel A. Changes in Log Relative Wages (%)													
High/Middle	19.4	21.2	19.1	13.0	0.0	-1.9	-7.3	20.5	20.2	19.7	2.8	2.9	4.3
Low/Middle	3.9	6.7	4.7	-1.4	-3.8	-5.7	-11.9	10.7	10.4	9.9	3.9	4.1	5.6
Panel B. Changes in Employment Shares (%)													
High	11.6	12.1	11.3	9.1	3.6	2.9	1.2	8.2	8.0	7.8	0.9	0.9	1.4
Middle	-14.4	-15.5	-14.4	-11.1	-6.1	-5.1	-2.2	-9.3	-9.1	-8.8	-1.6	-1.7	-2.5
Low	2.8	3.4	3.1	2.0	2.5	2.2	1.0	1.1	1.1	1.0	0.8	0.8	1.1
Panel C. Changes in Wage-Bill Shares (%)													
High	17.7	18.7	17.4	13.5	4.3	3.2	0.1	13.8	13.6	13.3	1.5	1.7	2.4
Middle	-19.2	-20.7	-19.1	-14.3	-5.9	-4.4	-0.2	-14.6	-14.4	-13.9	-2.3	-2.5	-3.5
Low	1.4	2.0	1.7	0.8	1.6	1.2	0.1	0.8	0.8	0.7	0.8	0.8	1.1
Panel D. Consumption-Index Moments													
	1980	2016											
Average C_t^h	0.069	0.064	0.109	0.134	–	0.106	0.129	–	0.072	0.085	–	0.064	0.074	–
Std. Dev. C_t^h	0.050	0.054	0.088	–	–	0.084	–	–	0.055	–	–	0.045	–	–
Panel E. Sectoral-Composition Metrics													
$100 \times \text{Cov}(\Delta sh_s, \alpha_H)$		0.440	0.671	0.483	-0.002	0.617	0.459	0.020	0.047	0.017	-0.024	0.240	0.256	0.349
$100 \times \text{Cov}(\Delta sh_s, \alpha_L)$		0.155	0.270	0.212	0.034	0.222	0.173	0.014	0.046	0.035	0.020	0.120	0.125	0.157

Notes: Panels A–C report 100 times the 1980–2016 changes in labor-market outcomes. Panel A entries are $100 \times \Delta \log(W_i/W_M)$ for $i = H, L$. Panels B and C report percentage-point changes. In each counterfactual block, the three columns report the baseline, representative-consumer, and homothetic versions of the model. Panel D reports the 1980 baseline value in the first numeric column and 2016 outcomes in the remaining columns. Dashes indicate not applicable entries.

across households.[31] These changes in the income distribution have no effect on labor-market outcomes in the homothetic model. With non-homothetic preferences, however, the resulting changes in the distribution of consumption indices shift value added from income-inelastic to income-elastic sectors. Because income-elastic sectors are relatively intensive in high- and low-skilled labor, this shift raises the relative demand for these occupations, increasing both their relative wages and their hours-worked shares. Hence, the income channel is the key reason why the combination of impulses has such different effects in the baseline and homothetic models.

Consumer heterogeneity strengthens the income channel because increases in the dispersion of expenditure across households shift demand towards income-elastic sectors (see Section 3.2). The contribution of income dispersion to the income channel can be evaluated by comparing the impact of the combination of impulses in the baseline and representative-consumer models. In the baseline, the covariances between the change in sectoral composition and the 1980 intensities in high- and low-skilled occupations are 0.67/100 and 0.27/100, respectively, whereas in the representative-consumer model they fall to 0.48/100 and 0.21/100. As a consequence of the smaller income channel, the relative wages of high- and low-skilled occupations increase by 2.1 and 2 log points less in the representative-consumer model than in the baseline, and their hours-worked shares increase by 0.8 p.p. and 0.3 p.p. less. Overall, consumer heterogeneity accounts for 25% of the income channel.

Roy impulses vs. factor intensity We next separate the labor-market and income distribution combination into two parts. The first corresponds to the change in factor intensities. The second corresponds to the simultaneous changes in the variances of productivity draws, the Pareto-tail parameter, and aggregate TFP. Table 6 shows that both combinations of impulses contribute to the income channel, although the contribution of changes in factor intensities is substantially smaller. For example, the differential effect between the baseline and homothetic models of the change in factor intensities on the relative wage of high-skilled workers is 0.8 log points, whereas the differential effect for the combination of changes in σ_{jt},κ_t and A_t is 7.3 log points. The same pattern holds for the contribution of the income channel to the relative wage of low-skilled workers and to the hours-worked shares of high and low-skilled workers. Therefore, although changes in factor intensities affect labor-market polarization through the income channel, most of the income channel is associated with changes in the variance of productivity draws, the Pareto tail, and aggregate TFP.

[31] In the baseline model, the average real consumption index across households increases from 0.069 in 1980 to 0.109 in 2016, and the standard deviation increases from 0.05 to 0.088 over the same period.

Changing A_{st} We already showed in Section 5.3 that indirect channels play a major role in accounting for the effects of changes in relative sectoral TFP on labor-market outcomes. We now assess whether these indirect effects operate through the income channel or the sectoral-composition channel by comparing the effect of this impulse across the three versions of the model. Table 6 shows that the impulse generates larger effects on the relative wages and hours-worked shares of high- and low-skilled workers in the homothetic model than in the baseline. This implies that the income channel does not contribute to the propagation of changes in relative sectoral productivity; all effects operate through the sectoral-composition channel.

Why does the income channel play a small negative role? Simply because changes in relative sectoral TFP slightly reduce both the mean and the dispersion of the distribution of consumption indices across households. In particular, the average consumption index decreases from 0.069 to 0.064, and the standard deviation falls from 0.05 to 0.045. The magnitude and direction of these changes give rise to a small, negative contribution of the income channel to polarization. Consequently, the impulse induces marginally greater polarization in the homothetic economy than in the non-homothetic versions of the model.

5.5 Quantification of the three channels

Next, we use simulations from Tables 5 and 6 to quantify the importance of the direct, sectoral composition and income channels in the evolution of labor-market outcomes.[32]

The direct channel contributes 13.0 log points to the high-skilled relative wage but lowers the low-skilled relative wage by 1.4 log points. The income channel adds 8.3 and 8.1 log points to the relative wages of high- and low-skilled workers, accounting for 42% of the high-skilled increase and fully offsetting the direct channel's negative effect for low-skilled workers. The sectoral-composition channel reduces relative wages by 2.1 and 3.0 log points. For hours-worked shares, the direct channel is the dominant force, increasing the high- and low-skilled shares by 9.1 and 2.0 percentage points. The income channel adds 3.0 and 1.4 percentage points, while the sectoral-composition channel subtracts 0.7 and 0.6 percentage points. Overall, the income channel explains 26% of the high-skilled increase and 51% of the low-skilled increase.

Turning to wage-bill shares, the direct channel raises the high-skilled share by 13.5 percentage points and the low-skilled share by 0.8 percentage points. The income channel adds

[32]The direct channel corresponds to the effect of the "Joint Labor/Expenditure impulses" in the homothetic version, Table 6. The income channel corresponds to the difference in the effects of the "Joint Labor/Expenditure impulses" in the baseline and the homothetic models (Table 6). The sectoral-composition channel corresponds to the "Joint relative sectoral TFP and taste parameters" in the baseline model (Table 5).

Table 7: Quantification of the Channels to Labor-Market Polarization

	Channel Decomposition				Comparison
	Direct Channel	Income Channel	Sect.-Comp. Channel	Total	Baseline Model All Impulses
Panel A. Changes in Log Relative Wages (%)					
High/Middle	13.0	8.3	-2.1	19.1	19.4
Low/Middle	-1.4	8.1	-3.0	3.7	3.9
Panel B. Changes in Hours-Worked Shares (%)					
High	9.1	3.0	-0.7	11.5	11.6
Middle	-11.1	-4.4	1.2	-14.3	-14.4
Low	2.0	1.4	-0.6	2.8	2.8
Panel C. Changes in Wage-Bill Shares (%)					
High	13.5	5.2	-1.2	17.6	17.7
Middle	-14.3	-6.4	1.7	-19.0	-19.2
Low	0.8	1.2	-0.6	1.4	1.4

Notes: "Direct Channel" reports the effect of simultaneously changing factor intensities, Roy productivity dispersion, aggregate TFP, and the Pareto tail of profit-income draws in the homothetic economy. "Income Channel" reports the additional effect of the same joint impulse in the baseline economy relative to the homothetic economy. "Sectoral-Composition Channel" denotes the indirect structural-transformation channel generated by the joint impulse to taste parameters and relative sectoral TFP. "Total" is the sum of the three channels. The last column reports the corresponding change in the calibrated baseline model when all impulses are fed simultaneously. Differences between "Total" and "Baseline Model All Impulses" reflect complementarities across the channel blocks. Panel A reports 100 $\times$ log changes in relative wages from 1980 to 2016. Panels B and C report percentage-point changes.

5.2 and 1.2 percentage points, and the sectoral-composition channel reduces them by 1.2 and 0.6 percentage points. In total, the income channel accounts for 29% of the rise in the high-skilled wage-bill share and 86% of the rise in the low-skilled share.

Taken together, these results show that the income channel is an important force behind labor-market polarization, far outweighing the sectoral-composition channel.

5.6 Robustness Exercises

We explore the robustness of our quantification by exploring six one-step departures from the baseline calibration. Across these exercises, the income channel remains sizable for relative wages and especially stable for hours worked.

Definition of the Income Channel We first consider an alternative definition of the income channel that includes the effect of relative sectoral TFP, A_{st}, on the expenditure distribution. Under this alternative definition, the income channel is the difference between the baseline and homothetic responses to the joint impulse composed of factor intensities,

Table 8: Income Channel Across Robustness Exercises

	Baseline	A_{st} in Income	Dispersion Target: Worker HH Inc.	Dispersion Target: Wage per Hour	No Pareto Tail	Correl. Roy	Nested NHCES
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Changes in Log Relative Wages (%)							
High/Middle	8.3	7.1	8.4	6.3	6.3	3.9	5.8
Low/Middle	8.1	6.4	8.2	6.0	6.3	8.9	3.9
Panel B. Changes in Hours-Worked Shares (%)							
High	3.0	2.6	3.2	3.3	2.3	3.9	2.3
Middle	-4.4	-3.7	-4.8	-4.8	-3.4	-5.0	-2.8
Low	1.4	1.1	1.5	1.5	1.1	1.2	0.5
Panel C. Changes in Wage-Bill Shares (%)							
High	5.2	4.5	5.5	5.1	4.0	5.1	3.9
Middle	-6.4	-5.3	-6.8	-6.3	-4.9	-6.2	-4.3
Low	1.2	0.9	1.3	1.2	0.9	1.2	0.3

Notes: The table reports the income channel under the baseline and six robustness exercises. The A_{st} column includes relative sectoral TFP in the income-channel block. The no-Pareto-tail column shuts down the profit-income draws Pareto tail in the baseline calibration. The worker HH Inc Roy target column replaces the baseline annual wage-income Roy variance target with worker-household income. The wage-per-hour Roy target column replaces the baseline Roy variance target with log hourly wage variance levels. The correlated Roy column reports the specification with three Roy correlation parameters. The nested NH-CES column uses a nested non-homothetic CES demand system with the baseline targets.

productivity draws dispersion, aggregate TFP, the Pareto tail for profit income, *and* relative sectoral TFP. Column (2) in Table 8 shows that the results are close to the baseline: the income channel raises high- and low-to-middle relative wages by 7.1 and 6.4 log points. The effects on hours worked are also similar, increasing the high- and low-skill hours shares by 2.6 and 1.1 p.p. and reducing the middle-skill share by 3.7 p.p.

Alternative Dispersion Measures to Calibrate the Occupational-Choice Block In the baseline calibration, the occupational-choice block targets the variance of log annual wage income. We consider two alternative targets to discipline the occupational-choice block: (i) the variance of log income among worker households, which is 0.632 in 1980 and 0.880 in 2016, and (ii) the variance of log hourly wages, which is 0.444 in 1980 and 0.545 in 2016. In both exercises, reported in columns (4) and (5) of Table 8, the income channel contributes significantly to the increase in relative wages and hours-worked shares of high- and low-skilled occupations. The hours responses are particularly stable across these alternatives, suggesting that the polarization effect of the income channel is not tied to a specific target of dispersion.

The Role of the Pareto Tail The previous exercises keep the baseline Pareto tail that governs profit-income draws fixed. We now ask how much of the income channel is accounted

for by the Pareto tail by shutting it down while leaving the rest of the calibration unchanged. Column (3) of Table 8 shows that the income channel remains very significant: relative wages rise by 6.3 log points at both ends of the skill distribution, and the hours-worked shares for high- and low-skilled workers increase by 2.3 and 1.1 p.p. Thus, the importance of the income channel does not hinge on the Pareto tail for profit income, but it certainly amplifies its magnitude.[33]

Allowing for Correlated Draws in the Occupational-Choice Block We next allow productivity draws in the occupational-choice block to be correlated across occupations, while keeping the demand system and the profit-income draws unchanged. To discipline the new covariance matrix of the occupational-choice block, in addition to relative wages per worker and wage bill shares, we target the total variance and within-occupation variance of log-hourly-wages.[34] The fitted correlations imply a strong negative correlation between low- and middle-skill productivity draws in both years, $\rho_{LM} = -0.93$ in 1980 and -0.91 in 2016, while the middle-high correlation rises from $\rho_{MH} = 0.11$ to 0.64. Under this calibration, the income channel remains sizable and produces hours reallocations close to the baseline, increasing high- and low-skill hour shares and reducing the middle-skill share. The contribution of household-heterogeneity to the income channel is also similar to the baseline, accounting for around 25% of it.

Richer Demand Substitution Patterns Finally, we replace the baseline non-homothetic CES demand system with a nested non-homothetic CES specification. We estimate an elasticity of substitution across three broad sectors (agriculture, services and non-service goods), and estimate separate within-nest elasticities for each broad sector. We find that all the estimated elasticities are less than one, indicating gross complementarity and, as expected, the elasticity of substitution across broad sectors is smaller than within broad sectors. Appendix B describes the estimation of the nested demand system and reports the estimated substitution elasticities. The income channel is somewhat smaller than in the baseline, but remains economically meaningful. It raises high/middle and low/middle relative wages by 5.8 and 3.9 log points, respectively, and continues to shift hours away from middle-skill occupations toward the high- and low-skill tails. The income channel continues to be the dominant chan-

[33] As expected, the contribution of household-heterogeneity to the income channel diminishes when the Pareto tail is shut down. It ranges from about 2.5% to 14% of the income channel across the different components of the baseline, and alternative-dispersion calibrations.

[34] Hourly wages are the natural target if the correlation across productivity draws arises from innate ability reflected in the number of efficiency units produced by the worker. The within-occupation hourly-wage variances are especially informative about the covariance matrix: changing correlations changes the set of workers selected into each occupation and therefore the dispersion of observed wages within occupations. See Appendix A.3 for further implementation details.

nel behind the rise of low/middle relative wages. The household-heterogeneity component also remains nontrivial, accounting for 23 percent of the income channel. Thus, the income channel remains substantive even when we relax the single substitution elasticity imposed in the baseline demand system.

5.7 Policy implications

So far, we have shown that distinguishing between the two indirect channels to labor-market polarization, sectoral composition and income, is essential because the income channel is far more important for the polarization of wages and hours-worked than the sectoral-composition channel.[35] However, another critical reason to differentiate the two indirect channels is that the policies that affect them are very different. To illustrate this, we provide examples of policies that influence each channel.

A wide range of policies affect labor-market outcomes through the sectoral-composition channel. Any price subsidy or tax directed to specific sectors alters relative prices and therefore the sectoral composition of the economy. Tariffs are a prominent example: by changing the prices of imported goods in a given sector—and in the sectors that use those goods as intermediates—they shift sectoral value-added shares. Innovation policies targeted at particular sectors operate similarly. Policies that raise innovation rates, adoption rates, or the intensity with which establishments use advanced technologies in specific sectors change relative sectoral TFP and thus the sectoral composition.[36]

Policies that affect the distribution of expenditures across households will instead impact the income channel. These include income transfers to low-income households, tax policies that modify after-tax income, and policies that influence savings rates differentially across the income distribution.

Therefore, distinguishing between the income and sectoral-composition channels is fundamental because the policies that influence each channel are entirely distinct.

6 Conclusion

We have developed a framework to study the polarization of labor markets. The framework features heterogeneous consumers with non-homothetic preferences and endogenous labor supply across three occupations à la Roy (1951). In addition to the direct effects that im-

[35]Two additional differences are that some of the impulses that drive the income channel (variance of productivity draws and aggregate TFP) exhibit complementarities with other shocks, and that the income channel depends on consumer heterogeneity.

[36]See Comin et al. (2025) and Cirera et al. (2026).

pulses have on relative labor demand and supply, our framework accommodates two indirect channels: one that operates through the sectoral composition of the economy, and one that operates through the distribution of income across households. The latter is the income channel. Because income-elastic sectors are intensive in high- and low-skilled occupations, impulses that increase either the average or the dispersion of household expenditure shift relative labor demand toward these occupations, thereby polarizing labor markets.

To quantify the income channel, we have developed two versions of the model that gradually shut it down: one that replaces heterogeneous consumers with a representative consumer and one where preferences are homothetic. The income channel plays a substantial role for the rise in relative wages for high- relative to middle-skilled workers and for the increase in hours-worked for high- and low-skilled workers, and it is the dominant factor in accounting for the rise of low-skill wages relative to middle-skill.

One quarter of the income channel operates through the effect of the shocks on the dispersion of household income, and the other three quarters through their impact on the average household income. These findings underscore the importance of non-homothetic preferences and consumer heterogeneity for studying macroeconomic phenomena, and labor-market polarization in particular.

Looking ahead, our framework can be extended in two promising directions. First, our analysis abstracts from savings and capital. Incorporating these elements would allow the model to capture the evolution of the labor share over time, as well as heterogeneity in savings rates and capital income across households. These additions would affect the income channel, although the net effect is, a priori, ambiguous.

Second, the income channel may be an important propagation mechanism for business-cycle shocks with distributional consequences. Pursuing this extension would require embedding our framework in a DSGE environment and developing solution methods capable of handling the absence of balanced growth that characterizes non-homothetic settings.

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.

Andreolli, M., N. Rickard, and P. Surico (2024): "Non-Essential Business Cycles," CEPR Discussion Paper 19773, CEPR Press, Paris and London.

Auclert, A. (2019): "Monetary Policy and the Redistribution Channel," *American Economic Review*, 109, 2333–2367.

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.

Blinder, A. S. (2009): "How Many US Jobs Might be Offshorable?" *World Economics*, 10, 41–78.

Blinder, A. S. and A. B. Krueger (2013): "Alternative Measures of Offshorability: A Survey Approach," *Journal of Labor Economics*, 31, 97–128.

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. I. Vizcaino (2021): "Skill-Biased Structural Change," *The Review of Economic Studies*, 89, 592–625.

Caron, J., T. Fally, and J. Markusen (2014): "International Trade Puzzles: A Solution Linking Production and Preferences," *The Quarterly Journal of Economics*, 129, 1501–1552.

——— (2020): "Per capita income and the demand for skills," *Journal of International Economics*, 123.

Cerina, F., A. Moro, and M. Petersen Rendall (2017): "The role of gender in employment polarization," LSE Online Documents on Economics 86170, LSE.

Cirera, X., D. Comin, and M. Cruz (2026): "Technology Sophistication Across Estab-

lishments," *The Quarterly Journal of Economics*, qjag018.

Comin, D., M. Dey, and M. Mestieri (2026): "Scale and Skill," Mimeo.

Comin, D., D. Lashkari, and M. Mestieri (2021): "Structural Change With Long-Run Income and Price Effects," *Econometrica*, 89, 311–374.

Comin, D., M. Mestieri, and A. Danieli (2020): "Income-Driven Labor Market Polarization," CEPR Discussion Papers 14980, C.E.P.R. Discussion Papers.

Comin, D. A., D. Lashkari, and M. Mestieri (2025): "The Structural Transformation of Innovation," Working Paper 33855, National Bureau of Economic Research.

Cozzi, G. and G. Impullitti (2016): "Globalization and Wage Polarization," *The Review of Economics and Statistics*, 98, 984–1000.

Cravino, J. and S. Sotelo (2019): "Trade-Induced Structural Change and the Skill Premium," *American Economic Journal: Macroeconomics*, 11, 289–326.

Dixit, A. K. and J. E. Stiglitz (1977): "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 67, 297–308.

Erosa, A., L. Fuster, G. Kambourov, and R. Rogerson (2025): "Labor Market Polarization and Inequality: A Roy Model Perspective," NBER Working Paper 33687, National Bureau of Economic Research.

Feenstra, R. C. and G. H. Hanson (1999): "The Impact of Outsourcing and High-Technology Capital on Wages: Estimates for the United States, 1979–1990," *The Quarterly Journal of Economics*, 114, 907–940.

Firpo, S., N. M. Fortin, and T. Lemieux (2011): "Occupational Tasks and Changes in the Wage Structure," IZA Discussion Papers 5542, IZA Network @ LISER.

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.

Grossman, G. M. and E. Rossi-Hansberg (2008): "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review*, 98, 1978–1997.

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.

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.

Jensen, J. B. and L. G. Kletzer (2010): "Measuring Tradable Services and the Task Content of Offshorable Services Jobs," in *Labor in the New Economy*, ed. by K. G. Abraham, J. R. Spletzer, and M. Harper, University of Chicago Press, chap. 8, 309–335.

Krueger, D. and F. Perri (2006): "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory1," *The Review of Economic Studies*, 73, 163–193.

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.

Murphy, D. (2016): "Welfare consequences of asymmetric growth," *Journal of Economic Behavior & Organization*, 126, 1–17.

Ngai, L. R. and O. Sevinc (2020): "A Multisector Perspective on Wage Stagnation," Working Paper 14855, CEPR.

Olivi, A., V. Sterk, and D. Xhani (2025): "Optimal Monetary Policy during a Cost-of-Living Crisis," Working paper, conditionally accepted, American Economic Review.

Orchard, J. (2025): "Non-homothetic Demand Shifts and Inflation Inequality," Finance and Economics Discussion Series 2025-085, Board of Governors of the Federal Reserve System.

Patterson, C. (2023): "The Matching Multiplier and the Amplification of Recessions," *American Economic Review*, 113, 982–1012.

Roy, A. D. (1951): "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 3, 135–146.

Werning, I. (2015): "Incomplete Markets and Aggregate Demand," NBER Working Papers 21448, National Bureau of Economic Research.

Online Appendix

A Data Description

In this section we briefly detail our data sources. We take it from previous work, and thus, we provide relatively brief descriptions and point the interested reader to the original papers for further details.

A.1 Labor-Market Outcomes

We follow Acemoglu and Autor (2011) in the construction of the baseline data on occupations, wages, and hours-worked shares. Here we provide a brief overview and refer the reader to the original work by Acemoglu and Autor for the details. The data for hours-worked comes from IPUMS USA and it includes the decennial censuses between 1980-2000 (with 10 years intervals) and annual data from the American Community Survey (ACS) between 2000-2016. The sample is restricted to individuals aged 16-64 who were employed in the previous year and are assigned to a known occupation (i.e., not missing or unemployed). We further restrict the sample to exclude the top and bottom 5% of the hourly wage distribution. Wage data comes from the Current Population Survey (CPS) to compute wages per occupation. We follow Acemoglu and Autor on this choice because the data in the ACS has only intervals starting in 2007. Occupations and industries are classified based on the 1990 Census Bureau classification scheme, which gives a consistent classification for all sample years. These industries are mapped to the BEA industry classification through mutual mapping to the NAICS codes. For each BEA industry, we compute the share of individuals within each occupation.

Our occupation classification is also taken from Acemoglu and Autor (2011). They divide the 382 original occupations into 4 broader categories that are characterized by their skill level: (1) managerial, professional and technical occupations; (2) sales, clerical and administrative support occupations; (3) production, craft, repair and operative occupations; and (4) service occupations. The first group is characterized by high-skill occupations, the second and third groups are characterized by middle-skill occupations and the last group is characterized by low-skill occupations. They measure skill by the average hourly wage of individuals in the occupation in 1980 where the mean wage in each occupation is calculated using workers' hours of annual labor supply times the Census sampling weights.[1]

[1]Our results on the negative correlation between occupation shares in middle-skill workers and income elasticity parameters are robust to decomposing middle-skill between groups (2) and (3).

A.1.1 Construction of Labor-Market Targeted Moments in the Calibration

Construction of wage-earnings dispersion moments. The overall and within-occupation wage-earnings dispersion moments are constructed from the IPUMS Census/ACS extract using annual wage income, `INCWAGE`. The extract is from IPUMS USA. We restrict the sample to individuals aged 16–64 with positive valid wage income, valid hours and weeks worked, and occupations that can be mapped into the Acemoglu-Autor occupation groups, as outlined above. We use person weights throughout.

For each year, we trim the bottom 5 percent of the weighted annual wage-income distribution. We then replace the upper tail above the 90th percentile with Pareto draws. The Pareto shape is chosen so that the corrected distribution matches the WID/Piketty-Saez top-10 share of wages and pensions.[2] The overall moment is the weighted variance of log corrected annual wage income. This gives 0.877 in 1980 and 0.968 in 2016.[3] Within-occupation moments are computed as the weighted variance of log corrected annual wage income within each occupation group, using the same corrected distribution.

For the robustness exercises in Section 5.6 we proceed as follows. For the exercise using worker income, we do an analogous procedure to described above, with the same sample restrictions, equivalization and top-correction to WID top-income shares, with the only difference that we use the worker-household HHINCOME series (instead of the annual wage income). For the exercise using the variance of log wage per hour, we compute the hourly wage obtained also from the ACS series trimmed at 1 percent.

Top-5 household-income target. The top-5 target used to calibrate the Pareto multiplier D_{ht} is constructed from CPS household income, `HHINCOME` from IPUMS CPS. The variables used are `HHINCOME`, `ASECWTH`, `AGE`, and `INCWAGE`. We keep worker households, defined as households with at least one member aged 16–64 with positive wage income, and weight households by `ASECWTH`. Household income is adjusted for household size by dividing `HHINCOME` by the square root of household size.

We then apply an analogous Pareto upper-tail correction. The CPS distribution below the 90th percentile is kept unchanged, while the upper tail is replaced by a Pareto tail calibrated to match tax-based top-10 pre-tax income shares from the Piketty-Saez/WID top-income series.[4] The resulting corrected CPS household-income distribution implies top-5 shares of

[2]The wage-tail targets come from the WID.world API, variable `sfiwag_p90p100_992_t`, the top-decile share of wages and pensions for the United States. The values used are 0.2687 in 1980 and 0.3676 in 2016. See WID.world, https://wid.world/, and the WID codes dictionary, https://wid.world/codes-dictionary/. The series was accessed on May 26, 2026.

[3]We have verified that replacing only the 95th percentile and repeating the procedure yields similar results.

[4]The targets used are 0.345 in 1980 and 0.475 in 2016. These are the Piketty-Saez/WID U.S. top-decile pre-tax income shares used in the top-correction file. See WID.world, https://wid.world/, and the updated

0.253 in 1980 and 0.378 in 2016, which are the targets used in our calibration.

A.2 Construction of Household Value-added Consumption Data

We start from household expenditure from the Consumer Expenditure Survey (CEX). We use data from the 2000-2001 period for our baseline results, and report estimates for 2000-2006 as a robustness check.[5] We follow the procedure described in Aguiar and Bils (2015) to clean the data. We restrict our sample to urban households. We drop households if they report spending less than 100 dollars in food per individual in the household over a three month time span, they have negative total or food consumption expenditure, total income is reported incomplete, they have not responded to all (four quarterly) interviews, income is below 50% of minimum wage, or if they earn money but do not work. To mitigate measurement error concerns, we drop the top and bottom 5% households according to their total income (after taxes) and we winsorize top and bottom 5% sectoral expenditures.[6] The only difference from Aguiar and Bils is that we keep all households with age of the reference person above 18. This allows us to capture the consumption of the elderly.

We then follow the procedure described in Buera et al. (2015) and convert the final good expenditures reported in the CEX into value-added expenditures using the BEA's 2000 input-output tables.[7] We do so by matching the finest level of expenditure categories in the CEX (called UCCs) to each sector in the BEA table. We start from the correspondence used in Buera et al. (2015), which takes the BLS crosswalk from UCC codes to PCE lines (supplemented with some expert judgement).[8] Then, as in Buera et al. (2015), we use the BEA crosswalk from PCE (table 245) to industries in the Input/Output table. For the few cases in which there are UCCs from 2000-2006 missing from their original list, we make the assignment to PCE lines based on our judgement. We attach the correspondence in the authors' websites. Following Comin et al. (2021), we also use sectoral, regional urban price series provided by the BLS.[9]

The Input-Output BEA sector codes in the Input-Output table that correspond to our

Piketty-Saez tables linked from Emmanuel Saez's data page, https://eml.berkeley.edu/~saez/; in particular, https://eml.berkeley.edu/~saez/TabFig2022.xlsx.

[5]We have experimented with different time frame periods between 1999 and 2007. The estimates are very stable across subsamples. Aguiar and Bils (2015) also report a similar finding in their estimated income elasticities.

[6]Total income after taxes is computed as in Aguiar and Bils (2015).

[7]We have also experimented using the detailed 2007 table, which contains over four hundred industries rather than sixty-nine. We obtain similar results. This exercise allows us to unpack industries such as "325 Chemical Products" into pharmaceuticals and the rest.

[8]This correspondence is available from the authors' website.

[9]We construct value-added price indices that are consistent with the value-added expenditure mapping; see the next section in this appendix.

groupings are:

- 6 for Education and Health Care, plus state and local expenditures in health and education. More specifically, we add the lines in the detail of the BEA value-added table "U.Value Added by Industry" "State and local government educational services" (line 185) and "State and local government hospitals and health services" (line 186). The BEA Table does not provide a breakdown for Federal expenditure in Education and Health, and we do not include it.
- 7 for Arts, Entertainment, Recreation and Food Services.
- G for Government (excluding the lines corresponding to state and local education and health care value added mentioned above).
- FIRE, PROF, 51, 81 for Finance, Professional, Information and other services (excluding gov't), excluding real estate (line 129, which includes housing and other real estate) in the BEA value-added table U.Value Added by Industry.
- 22, 23, 31G for Manufacturing, Mining and Utilities.
- 42, 44RT, 48T for Retail, Wholesale Trade and Transportation.
- 21 for Construction and Real Estate (line 129 in BEA Table U.Value Added by Industry).
- 11 for Agriculture.

A.2.1 Construction of Consumption Value-Added Price Indices

This subsection provides the construction of the price indices used in the demand estimation. Both the expenditure and price sides of the system are expressed in value-added terms, so that the demand system is internally consistent. We first state the value-added expenditure mapping and then derive the corresponding price index.

Value-added expenditures. Following the consumption value-added construction in Herrendorf et al. (2013), let e_h denote the vector of final expenditures of household h, let R denote the input-output total requirements matrix, and let v denote the vector of value-added shares in gross output. The value added embodied in household final expenditures is

$$x_h^{VA} = \langle v \rangle R \, e_h, \tag{A1}$$

where $\langle v \rangle$ denotes the diagonal matrix with vector v on the diagonal. The (s, i) element of $\langle v \rangle R$ measures the dollars of value added from sector s, direct and indirect, required to

deliver one dollar of final expenditure on good i. Equation (A1) defines the nominal value-added consumption object that we use as the dependent variable in the demand estimation.

Value-added price indices. An accounting identity for nominal expenditures does not, by itself, determine the price index for the constructed value-added composite. We therefore construct the price index associated with the same fixed-IO value-added consumption composite. Let

$$B \equiv \langle v \rangle R$$

denote the matrix that maps final expenditures into value-added expenditures. For value-added sector s, define the base-period IO weights

$$\omega_{si} = \frac{B_{si}\bar{e}_i}{\sum_k B_{sk}\bar{e}_k}, \tag{A2}$$

where $\bar{e}_i$ denotes aggregate final expenditure on category i in the base period. These weights measure the share of sector-s value-added expenditure that is embodied in final-expenditure category i.

We adopt a Cobb-Douglas assumption: the sector-s value-added consumption bundle is a fixed-share composite of the final-expenditure categories through which this value added is delivered,

$$C_s^{VA} = A_s \prod_i c_{si}^{\omega_{si}}, \qquad \sum_i \omega_{si} = 1,$$

where c_{si} is the component of final-expenditure category i associated with sector-s value-added consumption and A_s is a normalization. The dual unit-cost index for this composite is

$$p_s^{VA} = A_s^{-1} \prod_i \left(\frac{p_i^{FE}}{\omega_{si}} \right)^{\omega_{si}},$$

where p_i^{FE} is the price index for final-expenditure category i. Taking log changes eliminates the normalization constants and gives

$$\Delta \log p_s^{VA} = \sum_i \omega_{si} \, \Delta \log p_i^{FE}. \tag{A3}$$

This is a Stone price index for the fixed-IO value-added composite, with the IO expenditure shares in equation (A2) as weights. The same B that maps final expenditures into value-added expenditures (A1) also determines the weights that aggregate final-expenditure price indices into value-added price indices (A3).[10]

[10]If the fixed-IO composite were treated as Leontief, $C_s^{VA} = \min_i\{c_{si}/a_{si}\}$, the associated unit-cost index

A.2.2 Reduced-Form comparison with Aguiar and Bils (2015) and Imputations

Reduced-form expenditure elasticities We also estimate reduced-form expenditure elasticities following the specification in Aguiar and Bils (2015). Denoting a household by n, we estimate

$$\ln\left(\frac{x_{st}^n}{\bar{x}_{st}}\right) = \alpha_{str} + \eta_s \ln E_t^n + \Gamma_s Z^n + u_{st}^n, \tag{A4}$$

where x_{st}^n is value-added expenditure in sector s by household n during quarter t, $\bar{x}_{st}$ is average expenditure in sector s during quarter t, α_{str} denotes quarter-year-of-interview by region fixed effects, E_t^n is total quarterly expenditure, and Z^n includes demographic controls for age, number of earners, and household size.

We instrument total quarterly expenditure with total yearly after-tax income and income-quintile dummies, as in Aguiar and Bils (2015), to address measurement error in total expenditure. The estimated coefficients η_s provide reduced-form expenditure elasticities by sector and are reported in Column (4) of Table 3. They closely match the expenditure elasticities implied by our structural estimates, reported in Column (3).

Our implementation differs from Aguiar and Bils (2015) in three respects. First, we estimate elasticities over value-added consumption sectors rather than final-expenditure categories. Second, our mapping to health expenditures is different from Aguiar and Bils (2015). Third, because the value-added transformation and our eight-sector aggregation imply positive expenditures in each sector, we estimate the exact log specification in Equation (A4); Aguiar and Bils (2015) instead use a first-order approximation to accommodate zeros in detailed expenditure categories. We discuss these last two points in what follows.

Mapping of Health Expenditures Aguiar and Bils (2015) take the expenditure groupings for health and education from the CEX groupings. Instead, as we have discussed, we follow Buera et al. (2015) and map each expenditure category at the finest level of reporting in the CEX (called UCC) to different PCE lines and, ultimately, different NIPA lines. This does not make a difference for most of expenditure categories (e.g., education), but it makes a difference for health.

Health services in our data are composed of the categories that are mapped to lines starting with 62 in the BEA Input Output table (or lines 60 to 67 in NIPA table 2.4.5). These UCCs are: 340906 Care for elderly, invalids, handicapped, etc. (in the home), 560110 Physicians' services, 560210 Dental services, 560330 Lab tests, x-rays, 560400 Service by

would be linear in price levels, $p_s^{VA} = \sum_i a_{si}\, p_i^{FE}$. In log changes this corresponds to a weighted average with time-varying cost shares, $d\log p_s^{VA} = \sum_i \frac{a_{si} p_i^{FE}}{\sum_k a_{sk} p_k^{FE}}\, d\log p_i^{FE}$. The Cobb-Douglas/Stone approximation in equation (A3) is the fixed-share log-linear version of the same price-index problem. We have verified that our demand estimates go through under the Leontief specification.

professionals other than physicians, 570110 Hospital room (thru 2005 Q1), 570111 Hospital room and service (introduced 2005 Q2), 570210 Hospital service other than room (thru 2005 Q1), 570220 Care in convalescent or nursing home, 570230 Other medical care services, 570230 Other medical care services, 570240 Medical care incl. in homeowners expenses, 570901 Rental of medical equipment, 570903 Rental of supportive/convalescent equipment, 571230 Other medical care services, 572230 Other medical care services.

Instead, Aguiar and Bils follow the CEX grouping for "Health." This grouping contains UCCs from 540000 to 580902. This grouping includes expenditures in prescription drugs and medical supplies (UCC's 540000-550340, 570901, 570903) and Health Insurance (UCCs 580110-580902) as part of "Health." Instead, we match the UCC's corresponding prescription drugs to "Pharmaceutical and other medical products" (lines 40 and 41 in NIPA table 2.4.5) and health insurance to "Health Insurance" (line 93 in NIPA table 2.4.5). We are also including UCC "340906 Home health care" as a health service (which is not included in Aguiar and Bils).

Approximation of the log-ratio in the left-hand side of the Aguiar-Bils estimating equation Aguiar and Bils (2015) present their theory and justification for their estimating equation by having on the left-hand side of the regression $\ln(x_{hst}/\bar{x}_{st})$ (Equation 4 in their paper). However, in their empirical analysis, they substitute $\ln(x_{hst}/\bar{x}_{st})$ by its first-order approximation around $\bar{x}_{st}$, $(x_{hst} - \bar{x}_{st})/\bar{x}_{st}$. They justify their choice because the presence of zeros in the data (e.g., for education alone zeros account for around 50% of the observations). However, we do not have zeros in our data because it is (1) more aggregated (eight sectors rather than twenty) and (2) the input-output matrix makes it so that there is always some (albeit small) consumption of all eight industries (e.g., education and health is an input to other sectors and thus all households have positive value-added consumption of it).

Since we do not find a problem of zeroes arising in our value-added measure of consumption, we proceed with the estimation of the exact equation that has on the left-hand-side $\ln(x_{hst}/\bar{x}_{st})$. If we run our regression with the same approximation in the left-hand-side as Aguiar and Bils, we find similar coefficients. However, for the categories in which there is more dispersion in expenditures (which tend to be the more expenditure-elastic categories), this first-order approximation becomes worse and results in smaller estimates of the expenditure elasticity. This is especially true for Health and Education, in which we find that the coefficient can drop by almost 30%.

A.2.3 Demand Estimation: Robustness Checks

Our baseline estimation uses imputed consumption value-added for health, education and finance. In column (1) we report the estimated results if we add a college dummy in the

demand estimation for each sector indicating whether the household head has some college. The results remain very close to the baseline. Column (2) adds additional fixed effects to the baseline regression reported in column (0). In addition to sector-by-region fixed effects we interact them with year and round fixed effects. We obtain virtually the same estimates. In column (3) of Table A1 we report the estimates that we obtain without the imputation. The key difference is that the non-homotheticity parameter for Education and Health Care increases from 1.80 to 2.08. This implies an increase from 1.59 to 1.75 in the implied expenditure elasticity of the average household. The rest of the parameters are very similar. In particular, the imputation for Finance is not quantitatively important, with the parameter estimate changing from 1.39 in our baseline to 1.36. Columns (4) and (5) of Table A1 show that the estimated parameters are similar to our baseline if we remove year-round fixed effects and use the within year variation to also identify expenditure elasticities, or if we extend our sample period from 2000 to 2006.[11] Columns (6) and (7) report the reduced-form Aguiar-Bils specification with and without imputation. We find a very similar pattern as in columns (2) and (3).

Details on the Imputation Procedure We start discussing the imputation of education expenditures. We use expenditure per pupil at the school district-level for years 2000-2001 from the Common Core Database. For the median household income at the county level, we use the IPUMS-ACS. For counties with missing median household income, we impute the state median household income. We merge school districts to counties and regress log-expenditure per pupil on log household income with state fixed effects and a time trend. We then use the predicted values from the regression to impute the value of education of a student in K-12. We use the family files from the CEX to find out the number of children in each household in K-12 age. We impute the expenditure according to the number of children in K-12 age *if* the household does not report paying any elementary or high-school tuition. We impute it to UCC code 670210 "elementary and high-school tuition."

For Medicare and Medicaid expenditures we follow an analogous procedure. We use data in the Dartmouth Atlas on average expenditure per patient by hospital referral region (these data were generously shared and described to us by Douglas and Betsy Staiger and Jonathan Skinner). We then merge this information to county average household income. Since hospital referral regions do not coincide with counties, we use population in the referral region and county as a weight for computing the average expenditure per county. After this step, we proceed with the imputation in the same way as for education. Once we have a measure of expenditures per person in a given county, we use information in the CEX on the number of

[11]The reason why we start in 2000 and not earlier is that this is the start for disaggregated sectoral city price indices from the BLS.

Table A1: Estimated Demand and Income Elasticities: Robustness Checks

	Demand Parameters						Red. Form	
	(0)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Nonhomotheticity Parameters ε_s							*Inc. Elast.* ξ_s	
Education and Health Care	1.77	1.83	1.80	2.08	1.93	2.15	1.59	1.75
	(0.02)	(0.01)	(0.07)	(0.13)	(0.06)	(0.14)	(0.06)	(0.05)
Arts, Entertainment,	1.37	1.43	1.39	1.36	1.42	1.51	1.31	1.25
Recreation and Food Services	(0.03)	(0.01)	(0.03)	(0.04)	(0.03)	(0.06)	(0.03)	(0.03)
Finance, Professional, Information,	1.28	1.34	1.26	1.25	1.26	1.21	1.29	1.25
other services (excl. gov't)	(0.01)	(0.00)	(0.02)	(0.03)	(0.02)	(0.03)	(0.01)	(0.01)
Government[1]	1.00	1.04	1.00	1.00	1.00	1.00	1.08	1.07
	(-)	(0.00)	(-)	(-)	(-)	(-)	(0.01)	(0.01)
Manufacturing	0.74	0.77	0.69	0.71	0.61	0.67	0.91	0.92
	(0.03)	(0.00)	(0.03)	(0.03)	(0.03)	(0.04)	(0.01)	(0.01)
Retail, Wholesale Trade and	0.59	0.62	0.61	0.66	0.55	0.61	0.87	0.88
Transportation	(0.01)	(0.01)	(0.03)	(0.04)	(0.03)	(0.05)	(0.01)	(0.01)
Construction	0.48	0.51	0.48	0.49	0.44	0.47	0.80	0.80
	(0.01)	(0.01)	(0.04)	(0.06)	(0.04)	(0.06)	(0.02)	(0.02)
Agriculture	0.32	0.35	0.30	0.35	0.21	0.23	0.65	0.66
	(0.02)	(0.01)	(0.06)	(0.08)	(0.05)	(0.09)	(0.02)	(0.01)
Elasticity of Substitution σ	0.42	0.41	0.45	0.43	0.53	0.58		
	(0.02)	(0.01)	(0.05)	(0.08)	(0.04)	(0.06)		
College Control	N	Y	N	N	N	N	N	N
Region × Round FE	N	N	Y	Y	N	N	Y	Y
Region × Year FE	N	N	Y	Y	Y	Y	Y	Y
Region FE	Y	Y	Y	Y	Y	Y	Y	Y
Imputed Expenditures	Y	Y	Y	N	Y	Y	Y	N
Sample Years	00-01	00-01	00-01	00-01	00-01	00-06	00-01	00-01

Notes: Standard errors clustered at the household level in parentheses. Column (1) adds a college control to the baseline specification. [1]: The Government nonhomotheticity parameter is normalized to 1 in the baseline and robustness specifications; in column (1), the estimate is reported from the added-college-control specification.

household members under Medicaid and Medicare to make the imputation. In this case, we assign the expenditure to UCC 580901, which corresponds to "Medicare payments" (there is no analogous Medicaid payments in the CEX interview files until 2017).

Finally, we also explore the role of financial services that may potentially be underreported in the CEX. In particular, expenses in fund management that are subtracted from the fund payout appear to be missing in the CEX. For this reason, we proceed by assuming an expense ratio of 90 basis points over the year in all funds owned by a household and we evenly spread the expense over the 4 quarters (French, 2008).[12] These funds are pension funds (including

[12]The Investment Company Institute and Lipper report that the average expense ratios for bond funds in

Table A2: Within-Occupation Log Wage-Earnings Dispersion

	Model			Data		
	1980	2016	2016/1980	1980	2016	2016/1980
High	1.002	1.053	1.051	0.841	0.918	1.092
Middle	0.910	0.927	1.018	0.897	0.896	0.999
Low	0.897	0.911	1.015	0.935	0.885	0.946

Notes: The table reports within-occupation standard deviations of log wage earnings. Data moments use ACS annual wage earnings for workers aged 16–64 with positive valid wage income, valid hours and weeks, occupations mapped to Acemoglu-Autor groups, bottom-5-percent trimming, and a Pareto splice at the 90th percentile calibrated to tax-based top-10 wage-and-pension shares.

amount of money placed in a self-employed retirement plan) and the estimated market value of all stocks, bonds, mutual funds, and others such as securities. These expenses are imputed in UCC 710110 "Finance charges excluding mortgage and vehicle."

A.3 Correlated Occupational-Choice Model

We keep the demand system and the profit-income draws unchanged, and estimate the occupational-choice block allowing the three pairwise correlations, $(\rho_{LM}, \rho_{LH}, \rho_{MH})$, to vary freely subject to a positive-semidefinite correlation matrix. The calibration targets occupation wage-bill shares, high/middle and low/middle relative wages, the total variance of log hourly wages, and the within-occupation variances of log hourly wages for high-, middle-, and low-skill occupations.

The fitted correlations that we obtain are $(\rho_{LM}, \rho_{LH}, \rho_{MH}) = (-0.930, 0.262, 0.111)$ in 1980 and $(\rho_{LM},\ \rho_{LH},\ \rho_{MH}) = (-0.909, -0.265, 0.643)$ in 2016. As in the baseline, the calibration matches the non-variance labor-market moments closely. The within-occupation hourly-wage variance targets are also matched closely. In 1980, the targets for high-, middle-, and low-skill occupations are $(0.4039, 0.3940, 0.4330)$ and the model gives $(0.4046, 0.3965, 0.4334)$. In 2016, the corresponding targets are $(0.4808, 0.4396, 0.4312)$ and the model gives $(0.4831, 0.4443, 0.4318)$. The total log-hourly-wage variance is somewhat below the target: the model generates 0.422 in 1980 and 0.509 in 2016, compared with targets of 0.444 and 0.545. Since we are somewhat undershooting the overall variance, we view these results as generating a conservative value for the income channel.

2000 were 76 basis points, 89 for hybrid funds and 99 for equity funds. We take 90 which is a rough average between the three. In their study, they also document a substantial decline in these expense ratios over time. In 2015 they were 54, 77 and 68.

Table A3: Occupational-Choice Block and Expenditure-Tail Calibration Targets

	1980	2016
Variance Log Labor Income, target	0.877	0.968
Var($\log WB_{ht}$), model	0.877	0.965
Top-5 household-income share, target	0.253	0.378
Top-5 household-income share, model	0.273	0.378

Notes: Appendix A discusses the details of the construction of the targeted moments.

B Nested Non-homothetic CES preferences

As a robustness check, we re-estimate the demand system allowing for a nested non-homothetic CES structure. We partition the eight sectors of the baseline specification into three broad groups, $\mathcal{G} = \{a, m, s\}$, corresponding to agriculture, non-service goods, and services. Agriculture is a singleton. Non-service goods aggregates Construction and Real Estate (sector 23) with Manufacturing, Mining, and Utilities (sector 31G). Services aggregates Government (sector G), Arts, Entertainment, Recreation, and Food Services (sector 7), Health and Education (sector 6), Retail, Wholesale Trade, and Transportation (sector 30), and Finance, Professional, Information, and Other Services (sector 95).

At the upper nest, the household consumption index C_{ht} is defined by

$$\sum_{g \in \mathcal{G}} \left(C_{ht}^{\varepsilon_g} \zeta_{gt}\right)^{\frac{1}{\rho}} C_{hgt}^{\frac{\rho-1}{\rho}} = 1, \tag{B5}$$

where C_{hgt} denotes the consumption index of broad group g. For each non-singleton group $g \in \{m, s\}$, we define the lower nest as

$$\sum_{k \in \mathcal{S}_g} \left(C_{hgt}^{\varepsilon_k} \zeta_{kt}\right)^{\frac{1}{\rho_g}} c_{hkt}^{\frac{\rho_g-1}{\rho_g}} = 1, \tag{B6}$$

where $\mathcal{S}_m = \{23, 31G\}$ and $\mathcal{S}_s = \{G, 7, 6, 30, 95\}$. As in the baseline demand system, one taste shifter and one non-homotheticity parameter are normalized to one within each nest. We choose sector 31G as the reference sector in manufacturing, Government as the reference sector in services, and manufacturing as the reference sector in the upper nest.

The estimation follows exactly the logic in Section 4.2 and Appendix A.2. We use the structural equations implied by the demand system to estimate the relevant demand parameters. For each lower nest, let E_{hgt} denote expenditure in broad group g by household h at time t, let $x_{hkt} \equiv p_{kt} c_{hkt} / E_{hgt}$ denote the expenditure share of sector $k \in \mathcal{S}_g$, and let $\hat{k}_g$ denote

the reference sector in that nest. Then, for each $k \in \mathcal{S}_g \setminus \{\hat{k}_g\}$, the estimating equation is

$$\ln x_{hkt} = \Gamma_{kg} Z_h + \delta_{kgt} + (1-\rho_g) \ln \left(\frac{p_{kt}}{p_{\hat{k}_g t}} \right) + (1-\rho_g)(\varepsilon_k - 1) \ln \left(\frac{E_{hgt}}{p_{\hat{k}_g t}} \right) + \varepsilon_k \ln x_{h\hat{k}_g t} + u_{hkt}. \tag{B7}$$

Given the lower-nest estimates, we back out the corresponding taste shifters and construct the model-consistent lower-nest price indices,

$$P_{hgt} = \left[\sum_{k \in \mathcal{S}_g} \left(\zeta_{kt} p_{kt}^{1-\rho_g} \right)^{\theta_k} \left(x_{hkt} E_{hgt}^{1-\rho_g} \right)^{1-\theta_k} \right]^{\frac{1}{1-\rho_g}}, \tag{B8}$$

where $\theta_k \equiv (1-\rho_g)/\varepsilon_k$. We then estimate the upper nest with the same exact equation as in (20), replacing the sectoral prices in the baseline with the lower-nest price indices $\{P_{hat}, P_{hmt}, P_{hst}\}$ and replacing the eight-sector expenditure shares with expenditure shares over the three broad groups.

Results Table B4 shows that the nested specification leaves the main message of the baseline estimates unchanged. At the upper nest, agriculture remains strongly expenditure-inelastic, services remain expenditure-elastic, and the estimated elasticity of substitution is $\rho = 0.192$. Within the lower nests, the ranking of sectors based on their non-homothetic parameters within goods-producing and services is nearly identical to the implied rankings in the baseline eight-sector specification. In particular, manufacturing is more income-elastic than construction, and the ranking in services, from highest to lowest income elasticity, is: finance, health and education; arts and entertainment; government; retail and wholesale. The only difference with the baseline ranking is that finance now is more income-elastic than health and education and arts and entertainment. The elasticity of substitution within the lower nests is similar for services (0.423) and higher for goods (0.717). Taken together, the estimates are consistent with the baseline.

Quantification of channels To study the quantitative predictions of the model with the nested non-homothetic CES specification, we recalibrate the taste parameters in the baseline and homothetic versions of the model. Note that, because of Lemma 1, we do not need to recalibrate the remaining model parameters. We follow the same approach as in the baseline framework to compute the direct, income, and sectoral-composition channels.

Table B5 reports the effects of each channel on the labor-market outcomes. The direct and income channels remain the main forces behind the changes in relative wages, employment shares, and wage-bill shares, while the sectoral-composition channel is small for most out-

Table B4: Estimated Elasticities in the Nested Nonhomothetic CES

Nest / Sector	Elas. subst. ρ or ρ_g (1)	Income elast. param. ε_g or ε_k (2)
Panel A. Upper nest ($g \in \mathcal{G} = \{a, m, s\}$)		
Agriculture ($g = a$)		0.29
		(0.04)
Manufacturing ($g = m$)*		1.00
		(-)
Services ($g = s$)		1.63
		(0.04)
Elasticity of substitution	0.19	
	(0.03)	
Panel B. Non-service-goods lower nest ($k \in \mathcal{S}_m = \{23, 31G\}$)		
Construction ($k = 23$)		0.51
		(0.13)
Manufacturing, Mining, and Utilities ($k = 31G$)*		1.00
		(-)
Elasticity of substitution	0.72	
	(0.09)	
Panel C. Services lower nest ($k \in \mathcal{S}_s = \{G, 7, 6, 30, 95\}$)		
Health and Education ($k = 6$)		1.76
		(< 0.01)
Arts, Entertainment, Recreation, and Food Services ($k = 7$)		1.37
		(< 0.01)
Retail, Wholesale Trade, and Transportation ($k = 30$)		0.71
		(< 0.01)
Finance, Professional, Information, and Other Services ($k = 95$)		1.81
		(< 0.01)
Government ($k = G$)*		1.00
		(-)
Elasticity of substitution	0.42	
	(< 0.01)	

Notes: Standard errors clustered at the household level are shown in parentheses. Total number of households is 7,800. The estimation uses the same controls as the baseline demand regression. Column (1) reports ρ at the upper nest and ρ_g at lower nests. Column (2) reports ε_g in the upper nest and ε_k in the lower nest. *: The non-service-goods nest is the normalized reference sector in the upper nest, and sector 31G (Manufacturing, Mining and Utilities), and G (Government) are the normalized reference categories in the manufacturing and services lower nest, respectively.

comes. Relative to the baseline non-homothetic CES, the nested specification assigns a larger role to the direct channel and a smaller role to the income channel. Even so, the income channel remains economically meaningful. It raises high/middle and low/middle relative wages by 5.8 and 3.9 log points, raises high- and low-skilled employment shares by 2.3 and 0.5 percentage points, and raises high- and low-skilled wage-bill shares by 3.9 and 0.3 percentage points, continuing to shift hours away from middle-skill occupations toward the high- and low-skill tails. The household-heterogeneity component remains nontrivial, accounting for 23 percent

Table B5: Quantification of the Channels

	Channel Decomposition				All Impulses
	Direct Channel	Income Channel	Indirect ST Channel	Sum	Three Effects Simultaneously
Panel A. Changes in Log Relative Wages (%)					
High/Middle	12.9	5.8	0.1	18.8	19.4
Low/Middle	-1.6	3.9	0.6	3.0	3.9
Panel B. Changes in Employment Shares (%)					
High	9.1	2.3	0.0	11.4	11.6
Middle	-11.1	-2.8	-0.1	-14.0	-14.4
Low	2.0	0.5	0.2	2.6	2.8
Panel C. Changes in Wage-Bill Shares (%)					
High	13.5	3.9	0.0	17.4	17.7
Middle	-14.3	-4.3	-0.2	-18.7	-19.2
Low	0.8	0.3	0.2	1.3	1.5

Notes: The first four columns report the channel decomposition. "Direct Channel" reports the effect of simultaneously changing factor intensities, productivity-draw dispersion, the Pareto tail of profit-income draws, and aggregate TFP in the homothetic economy. "Income Channel" reports the additional effect of the same joint impulse in the baseline economy relative to the homothetic economy. "Indirect ST Channel" denotes the indirect sectoral-composition channel generated by the joint impulse to taste parameters and relative sectoral TFP. "Sum" is the sum of the three channels. The last column reports the corresponding change in the calibrated model when all impulses are fed simultaneously. Differences between "Sum" and "All Impulses" reflect complementarities across the channel blocks. Panel A reports 100 times log changes from 1980 to 2016. Panels B and C report percentage-point changes.

of the income channel. Moreover, the income channel remains the dominant channel for the evolution of the relative wage of low-skilled workers. Thus, relaxing the common substitution elasticity across sectors changes the allocation between the direct and income channels, but it does not overturn the main economic message.

C Aggregation with Heterogeneous Households

This appendix provides a second-order approximation that clarifies how cross-household expenditure dispersion affects aggregate sectoral demand under non-homothetic demand. We provide a general characterization and two specialized results. One under non-homothetic CES (NHCES) and another under the assumption of local isoelasticity of the expenditure elasticity and log-normal expenditures.

Let household h have total expenditure E_h, and define household expenditure on sector s as

$$g_s(E_h) \equiv p_s c_s(E_h).$$

Aggregate expenditure in sector s is

$$G_s \equiv \int_0^1 g_s(E_h)\, dh.$$

We compare this object with the expenditure generated by a representative household with the same mean expenditure,

$$\bar{E} \equiv \int_0^1 E_h\, dh,$$

which we denote $g_s(\bar{E})$.

Let

$$\xi_s(E) \equiv \frac{d \log g_s(E)}{d \log E}$$

denote the sectoral expenditure elasticity. Since $g_s'(E) = \xi_s(E) g_s(E)/E$, differentiating once more gives

$$g_s''(E) = \frac{g_s(E)}{E^2} \left[\xi_s(E)\big(\xi_s(E) - 1\big) + E\xi_s'(E) \right].$$

General result. A second-order Taylor expansion around $\bar{E}$ implies

$$\frac{G_s}{g_s(\bar{E})} \simeq 1 + \frac{1}{2} \left[\xi_s(\xi_s - 1) + \bar{E}\xi_s'(\bar{E}) \right] \frac{\mathrm{Var}(E_h)}{\bar{E}^2}.$$

where ξ_s is evaluated at $\bar{E}$. This expression is the general local aggregation result. The first term, $\xi_s(\xi_s - 1)$, is the curvature that would arise under a locally iso-elastic Engel curve. The second term, $\bar{E}\xi_s'(\bar{E})$, captures changes in the expenditure elasticity with household expenditure.

Proof. To obtain the approximation, write

$$G_s = \int_0^1 g_s(E_h)\, dh.$$

A second-order Taylor expansion of $g_s(E_h)$ around $\bar{E}$ gives

$$g_s(E_h) \simeq g_s(\bar{E}) + g_s'(\bar{E})(E_h - \bar{E}) + \frac{1}{2} g_s''(\bar{E})(E_h - \bar{E})^2.$$

Integrating over households, the first-order term vanishes because $\int_0^1 (E_h - \bar{E})\, dh = 0$, so

$$G_s \simeq g_s(\bar{E}) + \frac{1}{2} g_s''(\bar{E}) \mathrm{Var}(E_h).$$

Dividing by $g_s(\bar{E})$ and using

$$\frac{g_s''(\bar{E})\bar{E}^2}{g_s(\bar{E})} = \xi_s(\xi_s - 1) + \bar{E}\xi_s'(\bar{E})$$

yields the result.

C.1 A Simple Closed-Form Aggregation Benchmark

Assumption. Assume: (i) the Engel curve is locally iso-elastic around $\bar{E}$,

$$g_s(E) \simeq g_s(\bar{E}) \left(\frac{E}{\bar{E}}\right)^{\xi_s},$$

so that $\xi_s'(\bar{E}) = 0$; and (ii) household expenditure is log-normal, $\log E_h \sim \mathcal{N}(\mu, V)$, with $\bar{E} = \mathbb{E}[E_h]$.

Closed-form result. Under this assumption, the aggregation term has the closed form

$$\frac{G_s}{g_s(\bar{E})} \simeq \exp\left[\frac{1}{2}\xi_s(\xi_s - 1)V\right].$$

Discussion. This results offers a clean benchmark. The sign of the dispersion effect is governed by $\xi_s(\xi_s - 1)$. For normal necessities, $0 < \xi_s < 1$, expenditure dispersion lowers aggregate expenditure relative to a representative household with expenditure $\bar{E}$. For luxuries, $\xi_s > 1$, dispersion raises aggregate expenditure. Intuitively, the extra spending of high-expenditure households more than offsets the lower spending of low-expenditure households when the sectoral Engel curve is convex.

For two sectors s and r, the iso-elastic/log-normal approximation implies

$$\log \frac{G_s/g_s(\bar{E})}{G_r/g_r(\bar{E})} \simeq \frac{V}{2}\left[\xi_s(\xi_s - 1) - \xi_r(\xi_r - 1)\right].$$

Hence a mean-preserving increase in expenditure dispersion shifts aggregate demand toward sectors with larger $\xi_s(\xi_s - 1)$, in particular toward luxury sectors relative to normal necessities.

C.2 Aggregation under the NHCES demand system.

The term involving ξ_s' in the general result can be signed and bounded using the structure of preferences. Here we do this for NHCES. Let

$$x_j(E) \equiv \frac{g_j(E)}{E}, \qquad \bar{\varepsilon}(E) \equiv \sum_j x_j(E)\varepsilon_j, \qquad V_\varepsilon(E) \equiv \sum_j x_j(E)\big(\varepsilon_j - \bar{\varepsilon}(E)\big)^2.$$

In the NHCES demand system,

$$\xi_s(E) = \rho + (1-\rho)\frac{\varepsilon_s}{\bar{\varepsilon}(E)}.$$

Using $d\log x_j(E)/d\log E = \xi_j(E) - 1$ gives

$$\frac{d\bar{\varepsilon}(E)}{d\log E} = \frac{1-\rho}{\bar{\varepsilon}(E)} V_\varepsilon(E),$$

and therefore

$$E\xi_s'(E) = -(1-\rho)^2 \frac{\varepsilon_s}{\bar{\varepsilon}(E)^3} V_\varepsilon(E) \leq 0.$$

Thus, the correction term $E\xi_s'$ in the general result attenuates the static curvature $\xi_s(\xi_s - 1)$. Intuitively the term $E\xi_s'$ captures the departures from the isoelastic expenditure benchmark.

Combining the two terms in the general formula with their NHCES counterparts gives the equivalent expression

$$\xi_s(\xi_s - 1) + E\xi_s'(E) = \frac{1-\rho}{\bar{\varepsilon}^3}\left[\bar{\varepsilon}(\varepsilon_s - \bar{\varepsilon})\big(\rho\bar{\varepsilon} + (1-\rho)\varepsilon_s\big) - (1-\rho)\varepsilon_s V_\varepsilon\right],$$

where the right-hand side is evaluated at E. Hence local convexity in sector s is guaranteed if

$$\bar{\varepsilon}(\varepsilon_s - \bar{\varepsilon})\big(\rho\bar{\varepsilon} + (1-\rho)\varepsilon_s\big) > (1-\rho)\varepsilon_s V_\varepsilon.$$

This condition has a simple interpretation. The left-hand side is the static force toward convexity: it is positive only when $\varepsilon_s > \bar{\varepsilon}$, that is, when sector s has above-average non-homotheticity and hence an expenditure elasticity above one. The right-hand side is the attenuation coming from the fact that expenditure elasticities are not constant. As household expenditure rises, $\bar{\varepsilon}(E)$ rises because high-ε sectors receive larger expenditure shares, which lowers $\xi_s(E)$ for high-ε_s sectors. Thus luxuries are locally convex when their distance above the average, $\varepsilon_s - \bar{\varepsilon}$, is large enough relative to the cross-sector dispersion V_ε. If $\varepsilon_s < \bar{\varepsilon}$, the left-hand side is negative while the attenuation term is nonnegative, so the Engel curve is locally concave. The variance term can be bounded directly in terms of the support of the

non-homotheticity parameters. If $\varepsilon_{\min} \leq \varepsilon_j \leq \varepsilon_{\max}$, then

$$V_\varepsilon \leq (\varepsilon_{\max} - \bar{\varepsilon})(\bar{\varepsilon} - \varepsilon_{\min}).$$

This follows from $(\varepsilon_{\max} - \varepsilon_j)(\varepsilon_j - \varepsilon_{\min}) \geq 0$ and averaging with weights $x_j(E)$. The bound is sharper than the variance bound $(\varepsilon_{\max} - \varepsilon_{\min})^2/4$ because it uses the actual weighted mean $\bar{\varepsilon}$.

D Block-Recursive Calibration of Occupational-Choice, Profit-Income, and Demand Blocks

This appendix proves that the calibration is block-recursive in a precise sense. Taking as given the sectoral value-added shares and within-sector occupational wage-bill shares targeted in the calibration, firm optimality maps these objects into aggregate occupational wage-bill targets. Those targets, together with relative occupational wages and the log labor-income variance, identify the Roy block under Assumption 1. The profit-income tail is then calibrated from the resulting labor-income distribution and the top-income target. Only after these objects have been fixed does the demand system enter, through the sectoral taste shifters chosen to rationalize the same sectoral value-added shares.

Remark: calibration vs. counterfactual equilibrium. Lemma 1 is a statement about the calibration mapping conditional on the targeted production and labor-market moments. It does *not* imply that occupational sorting, profit-income draws, and demand are independent in counterfactual equilibria. In counterfactuals, changes in σ_t or in efficiency wages alter the labor-income distribution, and changes in the profit-income draws D_{ht} or their tail parameter κ_t alter the household expenditure distribution through Z_{ht}. With non-homothetic demand, these distributional changes affect sectoral demand and therefore feed back into equilibrium prices and allocations. The calibration is block-recursive because sectoral value-added shares are treated as moments to be matched. Thus, the result is not a claim that occupational choices and demand are independent in arbitrary counterfactual equilibria. It is a claim that, once the targeted production and labor-market moments are taken as given, the calibration equations for occupational sorting do not use demand parameters or profit-income draws.

Proof of Lemma 1 Fix a period t and write $VA_{st} \equiv P_{st}Y_{st}$. From firm optimality in (2), sectoral labor payments to occupation j satisfy

$$\widetilde{w}_{jt}\widetilde{X}_{jst} = \frac{\alpha_{jst}VA_{st}}{\psi_t}.$$

Hence the model-implied aggregate labor payment to occupation j is

$$WB^M_{jt} \equiv \sum_{s\in\mathcal{S}} \widetilde{w}_{jt}\widetilde{X}_{jst} = \sum_{s\in\mathcal{S}} \frac{\alpha_{jst}VA_{st}}{\psi_t}. \tag{D9}$$

Because the markup is common across sectors and occupations, it cancels from the occupational wage-bill shares:

$$\omega^M_{jt} \equiv \frac{WB^M_{jt}}{\sum_{k\in\{H,M,L\}} WB^M_{kt}} = \frac{\sum_{s\in\mathcal{S}} \alpha_{jst}VA_{st}}{\sum_{s\in\mathcal{S}} VA_{st}},$$

where we used $\sum_j \alpha_{jst} = 1$. Therefore, once the calibration takes the sectoral value-added shares and the occupation intensities $\{\alpha_{jst}\}_{j,s}$ as targets, the occupation wage-bill shares faced by the occupational-choice block are known before any demand object is computed. The aggregate level of labor payments fixes only a common wage scale.

Given wages per efficiency unit $\widetilde{\mathbf{w}}_t = (\widetilde{w}_{Ht}, \widetilde{w}_{Mt}, \widetilde{w}_{Lt})$, household h selects

$$o_{ht} = \underset{j\in\{H,M,L\}}{\arg\max} \{\eta_{hj}\widetilde{w}_{jt}\}.$$

Let $\mathcal{A}_{jt}(\widetilde{\mathbf{w}}_t)$ denote the corresponding selection set, as in (11), and let F_t denote the joint distribution of efficiency draws. The Roy block implies

$$\pi_{jt}(\widetilde{\mathbf{w}}_t, F_t) = \int_{\mathcal{A}_{jt}(\widetilde{\mathbf{w}}_t)} dF_t(\boldsymbol{\eta}), \qquad \widetilde{X}_{jt}(\widetilde{\mathbf{w}}_t, F_t) = \int_{\mathcal{A}_{jt}(\widetilde{\mathbf{w}}_t)} \eta_j \, dF_t(\boldsymbol{\eta}),$$

and therefore

$$WB^R_{jt}(\widetilde{\mathbf{w}}_t, F_t) = \widetilde{w}_{jt}\widetilde{X}_{jt}(\widetilde{\mathbf{w}}_t, F_t), \qquad W^R_{jt}(\widetilde{\mathbf{w}}_t, F_t) = \frac{WB^R_{jt}(\widetilde{\mathbf{w}}_t, F_t)}{\pi_{jt}(\widetilde{\mathbf{w}}_t, F_t)}.$$

where the superscript R is bookkeeping for "Roy" occupational-choice candidate, to distinguish these objects from the production-side wage-bill and wages targets. The relevant parameters of F_t in the baseline are the log-draw dispersions $\{\sigma_{jt}\}_{j\in\{H,M,L\}}$, with the means $\{\mu_{jt}\}_j$ normalized as in the main text. But, we notice that the same reasoning applies with correlated draws, as we do in Section 5.6.

The Roy equations are homogeneous in the vector of efficiency wages. For any $\lambda > 0$, the selection sets satisfy $\mathcal{A}_{jt}(\lambda \widetilde{\mathbf{w}}_t) = \mathcal{A}_{jt}(\widetilde{\mathbf{w}}_t)$. Thus occupation shares, efficiency-unit supplies, occupational wage-bill shares, relative average wages, and $\mathrm{Var}(\log WB_{ht})$ are unchanged by multiplying all efficiency wages by λ; only the level of labor income changes. It follows that the targeted Roy moments can be written as a system in the relative vector of efficiency wages and the draw dispersions:

$$\frac{WB^R_{jt}(\widetilde{\mathbf{w}}_t, F_t)}{\sum_{k\in\{H,M,L\}} WB^R_{kt}(\widetilde{\mathbf{w}}_t, F_t)} = \omega^M_{jt},$$

$$\frac{W^R_{jt}(\widetilde{\mathbf{w}}_t, F_t)}{W^R_{Mt}(\widetilde{\mathbf{w}}_t, F_t)} = \frac{W^{data}_{jt}}{W^{data}_{Mt}}, \qquad j \in \{H, L\},$$

together with the targeted variance of log labor income, $\mathrm{Var}(\log WB_{ht})$. If one instead writes the wage-bill conditions in levels, equation (D9) pins down the same common wage scale after the aggregate labor-payment level is chosen. This scale does not affect occupational choices, relative wages, occupation wage-bill shares, or the log labor-income variance. We proceed by assuming that this system of equations has a locally unique solution, so that it can be inverted.

Assumption 1 (Uniqueness) *For each period t, take as given the observed sectoral value-added shares and the sectoral occupation intensities $\{\alpha_{jst}\}_{j,s}$. The Roy moment system used in the calibration, namely occupation wage-bill shares, high/middle and low/middle average-wage ratios, and the log labor-income variance target, has a locally unique solution for the relative efficiency wages $\widetilde{\mathbf{w}}_t$ and the draw dispersions $\{\sigma_{jt}\}_j$, up to the common wage scale.*

By Assumption 1, these moment conditions determine $\{\sigma_{jt}\}_j$ and the relative efficiency wages $\widetilde{\mathbf{w}}_t$ using only the production and labor-market moments. No equation in this system contains the demand parameters ρ, $\{\varepsilon_s\}_{s\in\mathcal{S}}$, the taste parameters $\{\zeta_{st}\}_{s\in\mathcal{S}}$, household price indices P_{ht}, sectoral demand C_{st}, the cross-household demand aggregation rule, or the profit-income draws $D_{ht}(\kappa_t)$.

We now turn to household expenditures and incorporate the profit-income process to labor income. Labor income from the occupational choice model is

$$WB_{ht} = \max_{j\in\{H,M,L\}} \{\eta_{hj}\widetilde{w}_{jt}\}.$$

Profit income is given by (17): a share $1-p$ of households receives no profit income, while a share p receives $Z_{ht} = WB_{ht}D_{ht}$, with D_{ht} drawn from a Pareto distribution with tail

parameter κ_t. Hence final household expenditure can be written as

$$E_{ht} = WB_{ht} + Z_{ht} = WB_{ht}\left(1 + \mathbf{1}_{ht}^{Z} D_{ht}\right),$$

where $\mathbf{1}_{ht}^{Z}$ indicates that household h receives profit income. The draw $(\mathbf{1}_{ht}^{Z}, D_{ht})$ is independent of $\boldsymbol{\eta}_h$ and is realized after the occupational choice. It therefore changes the distribution of final household expenditure entering demand, but it does not enter the selection set $\mathcal{A}_{jt}(\widetilde{\mathbf{w}}_t)$, the efficiency-unit supplies $\widetilde{X}_{jt}$, occupation shares π_{jt}, occupation wage bills WB_{jt}, average wages W_{jt}, or the log labor-income variance targeted by the occupational-choice block.

The top-income target calibrates κ_t from the distribution of $E_{ht} = WB_{ht} + Z_{ht}$, taking the Roy labor-income distribution and p as given. This step also does not use the demand system. Moreover, the top-share condition is invariant to a common rescaling of WB_{ht}, so it does not feed back into the relative Roy wage vector or the draw dispersions. Thus the profit-income draws are an expenditure-layer object, not an occupational-choice object. Moreover, using the aggregate profit identity in the main text,

$$\int_0^1 Z_{ht}\, dh = p\, WB_t \frac{\kappa_t}{\kappa_t - 1} = (\psi_t - 1) WB_t,$$

then κ_t pins down the common markup ψ_t. This identity contains no demand parameters. Since ψ_t is common, it affects the aggregate split between labor payments and profits, and the common markup in the pricing equation, but not the occupational selection equations just described.

After the occupational-choice and profit-income objects have been determined, the demand block is used to choose the taste parameters that rationalize the same sectoral value-added shares. The sectoral price equation is

$$P_{st} = \psi_t \frac{\prod_j \left(\frac{\widetilde{w}_{jt}}{\alpha_{jst}}\right)^{\alpha_{jst}}}{\tilde{A}_{st}},$$

and aggregate demand in the heterogeneous-household non-homothetic CES economy is

$$C_{st} = \zeta_{st} \int_0^1 C_{ht}^{\rho+\varepsilon_s} \left(\frac{P_{ht}}{P_{st}}\right)^{\rho} dh.$$

The taste parameters $\{\zeta_{st}\}_s$ are then chosen so that the model matches the targeted sectoral value-added shares. Replacing this demand block with the representative-consumer or homothetic version changes the taste parameters required to rationalize those shares. It does not change the occupational-choice parameters $\{\sigma_{jt}\}_j$ and $\widetilde{\mathbf{w}}_t$, nor the profit-tail parameter κ_t,

because none of those calibration equations uses the demand system.

This proves the block-recursive calibration structure stated in the lemma.