The Luxury Shift: Productivity Drag and Distributional Implications

Joern Onken *

University College London

Job Market Paper

November 17, 2026

[Please click here for the latest version]

Abstract

Sectoral reallocation has been a drag on U.S. productivity growth. This paper shows that the expansion of luxury services accounts for most of this drag. Yet the forces behind their expansion disproportionately benefit poorer households. Linking household expenditure data to producing industries, I distinguish luxuries from necessities *within* both goods and services. Luxury services are labor-intensive, have weak productivity growth and disproportionately employ low-paid workers. Their expansion accounts for about two-thirds of the 0.20-percentage-point annual drag on labor productivity growth over 1980–2024. To explain this reallocation and its distributional consequences, I estimate a nonhomothetic demand system and embed it in a four-sector general equilibrium model with heterogeneous households. Rising incomes shift spending toward luxury services, supporting labor earnings and lowering capital returns. Slow productivity growth raises the relative price of luxury services, which further increases their expenditure share because substitution to other sectors is limited. Poorer households benefit from stronger demand for low-paid labor and are less exposed to lower capital returns and rising luxury-service prices. Relative to an economy with homothetic demand and equal sectoral productivity growth, average welfare gains in 2024 equal 1.9% of initial budgets for the bottom expenditure decile, versus losses of 2.6% for the top.

Keywords: Structural transformation, nonhomothetic preferences, consumption, productivity growth, inequality, Baumol's cost disease.

JEL Classification: D12, D31, E21, E25, O41, O47.

*Email: joern.onken.20@ucl.ac.uk. I am indebted to my advisors, Vincent Sterk and Morten Ravn, for their invaluable guidance and support throughout this project, and to Lukasz Rachel and Marta Morazzoni for countless very helpful discussions. I am grateful to Javier Boncompte, Timo Boppart, Adrien Couturier, Wei Cui, Ararat Gocmen, Charles Gottlieb, Elisa Keller, Hannes Malmberg, Johannes Matt, Matthew Nibloe, Lukas Nord, Alan Olivi, Federico Pessina, Franck Portier, Hugo Reichardt, Kjetil Storesletten, Ludwig Straub and Davide Zufacchi, as well as conference and seminar participants at the UCL-Sciences Po Macro Workshop, NorMac, the EEA and UCL for their comments and feedback.

1 Introduction

As households become richer, they spend a larger share of their budgets on luxuries such as eating out, education, consumer electronics, travel and recreation.[1] This luxury shift is one of the most predictable features of economic growth. Its macroeconomic consequences, however, depend on which luxuries expand and how they are produced. Some luxuries are produced in industries with rapid productivity growth, while others are labor-intensive and have much weaker productivity growth. Growing demand for the latter could slow aggregate productivity growth while increasing demand for the workers they employ, including lower-paid workers. This raises two questions: Does economic growth slow itself down as rising incomes shift resources toward luxury sectors with weak productivity growth? And who gains from the forces behind this shift, given that the households buying luxuries are often not the ones producing them?

This paper shows that luxury sectors account for almost all of the U.S. productivity drag from sectoral reallocation over the last decades. Over 1980–2024, observed U.S. labor productivity growth is 0.20 percentage points per year lower than without sectoral reallocation. Using a new demand-based industry classification that distinguishes luxuries from necessities within both goods and services, I find that luxury sectors contribute 0.19 percentage points to this gap. The two luxury sectors have moved in opposite directions: Luxury services, a sector with weak productivity growth, have expanded over time and account for about two-thirds of the drag on their own. Luxury goods, a sector with strong productivity growth, have contracted in relative terms and account for most of the rest.

To uncover the forces driving this reallocation and its distributional implications, I develop a macroeconomic model with household heterogeneity and an estimated, nonhomothetic demand system. I find that the expansion of luxury services is driven by two main forces. First, overall income growth expanded luxury sectors, as their income elasticities are the highest among all sectors. Second, the price of luxury services increased significantly relative to that of luxury goods. This further raised the expenditure share of luxury services due to limited substitutability between luxury goods and services. Surprisingly, a significant part of these effects originates in strong productivity growth in the luxury goods sector itself. This growth raised the relative price of luxury services and boosted overall income growth. The sector with the highest productivity growth thus helped shrink its own weight in the economy.

The forces behind the expansion of luxury services also affect households differently. Households differ in their income sources, consumption baskets and occupations, and thus in their exposure to sectoral change through consumption prices, wages and returns on capital. I find that these forces disproportionately benefited poorer households. They spend smaller budget shares on these services, so they were less exposed to the rise in their relative price. They also work relatively often in luxury services, so they benefited from the higher labor demand in this sector. Richer households, by contrast, rely more on income from capital, whose returns fell as demand shifted toward labor-intensive luxury services.

[1] These categories are luxuries in the economic sense: Their budget shares rise with total household expenditure at given prices (Deaton and Muellbauer 1980; Aguiar and Bils 2015).

The starting point of this paper, and one of its key contributions, is a new way of classifying production. The structural transformation literature typically divides the economy into goods and services, sometimes further splitting services according to production characteristics such as productivity growth or skill intensity. However, these production-based classifications do not tell us which industries attract a growing share of household spending as incomes rise. I therefore classify industries from the demand side. To do so, I estimate Engel curves for 35 consumption categories in the Consumer Expenditure Survey (CEX), following Aguiar and Bils (2015), and use the resulting expenditure elasticities to distinguish luxuries from necessities. I then map these categories to detailed producing industries and use input-output links to extend the classification to their upstream suppliers. Combining the resulting luxury/necessity labels with the goods/services distinction based on industry codes yields four sectors: Necessity goods (food manufacturing), necessity services (utilities), luxury goods (furniture manufacturing) and luxury services (restaurants). The classification covers 297 four-digit NAICS industries (accounting for 99.1% of private-sector employment in 2023) and is constructed independently of their productivity growth and factor intensity. Differences in production characteristics between the four sectors are thus findings of the paper, rather than features imposed by the classification.

My classification reveals important differences in productivity growth, consumer-price growth, labor shares and employment patterns that a conventional goods/services split conceals. Within services, employment has increasingly shifted toward luxury services, which have weaker labor productivity growth and a higher labor share than necessity services. Luxury services also disproportionately employ lower-paid workers within both college and non-college groups, which suggests that their expansion can increase demand for these workers' labor. The contrast between the two luxury sectors is equally striking: Luxury goods have experienced falling consumer prices and contracting domestic employment, while luxury services have experienced the fastest growth in both consumer prices and employment among the four sectors.[2]

I quantify the drag on aggregate labor productivity growth from sectoral reallocation (Baumol 1967; Duernecker, Herrendorf, and Valentinyi 2024) using growth accounting. I compare observed growth in real value added per hour in U.S. private-sector industries over 1980–2024 with a benchmark that fixes initial nominal value-added weights across the four sectors and excludes labor reallocation. Both calculations retain the observed sectoral productivity paths, including changes in industry composition within each sector. The accounting shows that luxury goods and luxury services contribute to the productivity drag through different channels. For luxury services, most of the contribution comes from hours shifting toward these services, where value added per hour is below average. For luxury goods, it comes mainly from their declining share of nominal value added, which reduces the weight on their relatively strong productivity growth. This growth accounting exercise does not, however, explain which changes in demand or technology drove the sectoral reallocation.

[2]The main patterns in prices, productivity, employment, and labor shares across the four sectors are robust to excluding housing and healthcare from the classified categories and to using an alternative Engel specification. See Appendix B.1.

I therefore examine next *why* luxury services expand and find that both rising household incomes and relative-price changes contribute substantially. I quantify their respective roles using a nested nonhomothetic demand system estimated on U.S. CEX consumption microdata, building on Comin, Lashkari, and Mestieri (2021). I measure consumption as the domestic value added generated by household purchases, including the value added of upstream suppliers. I estimate income effects by comparing how households with different expenditure levels spend their budgets, and substitution elasticities using variation in relative prices across regions and over time. I then use the estimated demand system, together with observed household budgets, prices, and demographic and regional characteristics, to predict how sectoral expenditure shares have changed over time. I calibrate taste intercepts to match initial sectoral expenditure shares in 1999, then hold these intercepts and the estimated demand parameters fixed thereafter. I also exclude the estimated calendar effects from the predictions. Over 1999–2023, the demand system predicts a 2.3-percentage-point increase in the luxury-service share, compared with 2.8 points in the data. Of the predicted increase, compensated relative-price effects contribute 1.3 percentage points, compared with 0.8 points from income effects.

The large contribution of relative-price changes reflects limited substitution away from luxury services: As these services become relatively more expensive over time, households cut their consumption less than proportionately, which raises the luxury-service expenditure share. Almost all of the income-effect contribution comes from growth in average household expenditure, rather than changes in expenditure inequality. This implies that economic growth itself, by increasing households' purchasing power, contributes to reallocation toward luxury services with weak productivity growth. Future income growth can therefore sustain this shift and prolong its productivity drag, provided productivity growth in luxury services remains weaker than elsewhere.

The historical demand decomposition takes household budgets and prices as given. To make these budgets and prices endogenous, I therefore embed the estimated preferences in a four-sector general equilibrium model with heterogeneous households. Sectors differ in productivity growth, capital intensity, and the types of labor they employ, while households differ in their consumption baskets and sources of income. Workers retain their occupation type but can move across sectors. Wages, capital returns, and sectoral prices adjust to clear markets. I calibrate the model using sectoral production data and microdata on household expenditure, labor earnings, and wealth. I study a sequence of static equilibria along historical paths for technology and aggregate factor supplies. This allows me to assess how changes in demand and sectoral productivity affect sectoral reallocation, factor returns, and household purchasing power.

In the model, nonhomothetic demand contributes more to the expansion of luxury services than uneven sectoral productivity growth. Over 1997–2024, the luxury-service share of nominal value added rises by 7.5 percentage points in the baseline, compared with 2.4 points under homothetic demand and 5.8 points under equalized sectoral productivity growth. A third counterfactual combines both changes. The homothetic counterfactual preserves each household's

initial consumption basket and compensated price responses. All economies share the same initial allocation, household factor ownership, and paths for aggregate capital and labor supply.

Who gains from the forces behind the luxury shift, and who loses? I evaluate each household's consumption opportunities across the model economies using its initial nonhomothetic preferences. I express welfare differences as differences in equivalent expenditure at common initial prices, relative to initial budgets. In 2024, relative to an economy with homothetic demand and common sectoral productivity growth, the baseline yields average welfare gains of 1.9% of initial budgets for the bottom decile of initial size-adjusted expenditure, compared with losses of 2.6% for the top. These welfare differences reflect changes in both household incomes and consumer prices. Relative to the homothetic benchmark, nonhomothetic demand raises real wages most for the lowest- and highest-paid occupation groups, both used intensively by luxury services, while lowering capital returns. For households in the bottom decile, higher labor income more than offsets losses from lower capital returns on average, whereas for those in the top decile it does not. Poorer households also spend smaller shares of their budgets on luxury services and are therefore less exposed to increases in their relative price. Uneven productivity growth likewise generates average gains at the bottom and losses at the top relative to the common-growth benchmark, mainly through consumer prices. The average distributional pattern survives alternative financing rules, occupation groupings, labor-substitution elasticities, and constant rather than time-varying occupation-demand weights (Appendices E.6, E.7, and E.8). However, these averages conceal substantial heterogeneity: Every expenditure decile contains households that gain and households that lose. These counterfactuals change demand and sectoral productivity growth, not just sectoral allocations. The welfare differences therefore are not due to reallocation alone, nor do they measure households' overall gains from historical growth. What they show is that the forces driving the expansion of luxury services disproportionately benefit poorer households, even as the associated reallocation reduces aggregate productivity growth.

Finally, I examine how productivity growth in each sector contributes to the expansion of luxury services. I use a Shapley decomposition to separate the contributions of nonhomothetic demand and each sector's productivity growth relative to a common reference-growth benchmark. Nonhomothetic demand accounts for about 69% of the model's baseline increase in the luxury-service share of nominal value added. Relatively strong productivity growth in luxury goods accounts for 13%, and relatively weak growth in luxury services for 10%; the contributions from the necessity sectors are small. These results reveal an important connection between the two luxury sectors: Rapid productivity growth in luxury goods helps expand luxury services by raising household purchasing power and making these services more expensive relative to luxury goods, even though growth accounting attributes most of the productivity drag to luxury services. This connection would be obscured by grouping all luxuries or all services together.

1.1 Related Literature

Structural change and productivity. The idea that labor can shift toward sectors with slower productivity growth goes back to Baumol's cost disease argument (Baumol 1967). Ngai and Pissarides (2007) formalize this mechanism, showing that differences in sectoral productivity growth shift labor toward slower-growing sectors when consumption goods are complements. Closest to my analysis, Duernecker et al. (2024) distinguish progressive from stagnant services based on their productivity growth. They find these two sectors to be substitutes, which means households substitute away from stagnant services as they become more expensive over time. This limits the drag of stagnant services on future productivity growth, even though these services are luxuries.

In contrast, rather than classifying industries by their productivity growth or factor intensity, I start from demand. I use household spending patterns to distinguish luxuries from necessities, and then link these goods and services to the industries that produce them. Only then do I compare the sectors' productivity growth and factor intensity. This matters because I want to find out whether sectors classified from household demand differ in how they produce, rather than assuming these differences from the start by construction. Furthermore, my estimated demand system implies that necessity and luxury services are gross complements, and that limited substitution away from luxury services allows their expenditure share to rise as they become relatively more expensive.[3]

The demand analysis builds on Foellmi and Zweimüller (2008), Boppart (2014) and Comin, Lashkari, et al. (2021), who show how nonhomothetic preferences shape structural change.[4] I estimate a nested four-sector version of the nonhomothetic CES demand system developed by Comin, Lashkari, et al. (2021). This demand system allows me to study the shift to luxuries separately from the shift to services. Through the nesting, I allow for separate substitution elasticities between luxuries and necessities, as well as between goods and services within each group. Herrendorf, Rogerson, and Valentinyi (2013) show that the relative importance of income and price effects depends on whether consumption is measured as final expenditure or embodied value added. I therefore estimate my demand system using the domestic value added generated by household purchases, including that of upstream suppliers.

Income effects and labor demand. Comin, Danieli, and Mestieri (2026) show that rising incomes shift spending toward sectors that rely heavily on low- and high-skill occupations, which generates labor-market polarization. This income-driven labor-demand mechanism is also part of my model: As rising incomes increase demand for luxury services, wages rise most for the occupation groups that are used most intensively in that sector. Related work links spending by more-skilled households to demand for less-skilled labor (Mazzolari and Ragusa 2013), trade-

[3] Ho, Lashkari, and Sarte (2026) study how investment and intermediate-input linkages transmit sectoral productivity shocks to future GDP. My model instead takes aggregate factor-supply paths as given and focuses on how nonhomothetic demand reallocates production and affects household welfare.

[4] Foellmi and Zweimüller (2006) study how nonhomothetic demand and the income distribution affect innovation incentives. Comin, Lashkari, and Mestieri (2025) study how these demand shifts redirect innovation across sectors and shape productivity growth. I instead take sectoral productivity paths as given and quantify how demand reallocates production across existing industries.

induced relative-price changes to the skill premium (Cravino and Sotelo 2019), and income-driven expenditure shifts to the aggregate labor share (Hubmer 2023). Wiedemann (2026) links inequality to structural change, while Ngai and Sevinc (2025) show how uneven productivity growth can dampen the wages of low-skilled workers. Andreolli, Rickard, and Surico (2025) document that non-essential industries disproportionately employ low-paid and hand-to-mouth workers, whose earnings are particularly exposed to contractions in non-essential spending.

My paper builds on this income-driven labor-demand channel but connects it to the productivity consequences of structural change and to household welfare. My four-sector classification reveals that luxury services combine weak productivity growth with intensive use of both low- and high-paid occupations. Their expansion therefore contributes to the measured productivity drag from sectoral reallocation while increasing demand for these occupation groups. I then quantify how nonhomothetic demand and uneven productivity growth affect household welfare through occupational wages, capital returns, and consumer prices. Crucially, occupation rank does not directly correspond to household expenditure rank: Households differ in the types of labor they supply, their capital ownership, and their consumption baskets. High-paid occupation groups can therefore gain even as high-expenditure households lose on average. The occupational wage response alone does not determine household welfare. Thus, while Comin, Danieli, et al. (2026) focus on occupational polarization, I connect the same income-driven demand mechanism to the productivity drag from sectoral reallocation and to the distribution of welfare across households.

Structural change and household welfare. Baqaee and Burstein (2023) show that measured output growth and welfare can diverge when structural change is driven by income effects or taste changes. Jaravel and Lashkari (2024) develop methods to measure growth in consumer welfare with income-dependent preferences. Bergman, Jaimovich, and Saporta-Eksten (2026) show that the distributional welfare effects of sectoral changes depend on both consumer prices and relative wages, while Olivi, Sterk, and Xhani (2026) study how nonhomothetic demand and differences in households' consumption baskets shape the effects of sectoral shocks and monetary-policy trade-offs. Fan, Peters, and Zilibotti (2023) quantify the gains from productivity growth in nontradable consumer services in India, finding that richer urban households benefit disproportionately.

My paper brings this welfare perspective to the long-run shift toward luxury services. I quantify how nonhomothetic demand and uneven sectoral productivity growth affect household welfare through labor earnings, capital returns, and consumer prices in general equilibrium.[5] The measured productivity drag from sectoral reallocation does not by itself reveal who gains from the demand and productivity changes behind it.

The remainder of the paper is organized as follows. Section 2 introduces the industry classification, documents sectoral differences in production and household exposure, and quantifies the productivity drag from sectoral reallocation. Section 3 presents and estimates the demand system and decomposes historical changes in sectoral expenditure shares. Section 4 develops the general equilibrium model, describes its calibration, and assesses its fit. Section 5 presents the

[5] In related work on the distributional implications of technological progress, Moll, Rachel, and Restrepo (2022) show that automation can increase inequality by raising returns to wealth.

counterfactuals, their household welfare effects, and the contributions of sectoral productivity growth to luxury-service expansion. Section 6 concludes.

2 The luxury shift in the data

This section documents the empirical patterns underlying the luxury shift. After describing the data and the four-sector classification, I show that luxury goods and luxury services follow very different productivity and employment paths: Domestic employment contracts in luxury goods, where productivity grows rapidly, and expands in luxury services, where productivity growth is weak. Luxury services also have the highest labor share, linking their expansion more closely to labor earnings. I then show that higher-income households spend larger shares of their budgets on these services, while these services account for a larger share of hours worked by lower-earning workers. These patterns motivate the distributional welfare analysis later in the model, because they highlight that households are affected by the luxury shift (sectoral reallocation to luxury services) both through what they consume and through how they earn their income. Finally, growth accounting shows that luxury services account for most of the productivity drag from sectoral reallocation.

2.1 Data and the classification of luxuries

I classify industries as luxuries or necessities using the expenditure elasticities of the consumption they produce. Luxuries have elasticities greater than one, meaning their budget shares rise with total household expenditure. Necessities have elasticities below one (budget shares decreasing in total expenditure). Hence, my luxury–necessity distinction is based on household expenditure responses (demand), rather than on the productivity or employment patterns examined below (which are production-related).

I estimate Engel curves for 35 granular consumption categories using the Consumer Expenditure Survey (CEX) over 1980–2020. For each category, I regress household spending relative to the category-year mean on log total household expenditure, controlling for household characteristics and year effects. To address measurement error in reported expenditure, I use the two instrumental-variable specifications of Aguiar and Bils (2015). In Specification I, which I use as my baseline, I instrument log total expenditure with income-group indicators and log after-tax household income. In Specification II, I instead only use spending from the final two expenditure interviews and instrument log total expenditure with its counterpart from the first two interviews. Appendix A.1 provides the full specification and sample details, and Appendix Table 7 reports the estimates for all 35 categories. I use the pooled Engel estimates to classify industries and keep this classification fixed over time. This way, changes in sector shares reflect actual reallocation of economic activity, rather than industries switching between luxuries and necessities. For 28 of the 35 consumption categories, the point-estimated expenditure elasticities are on the same side of one in 1984–1986 and 2017–2019. Appendix Table 8 reports the estimates and the categories that change status.

I map these 35 categories to four-digit NAICS producing industries, for example the CEX consumption category "Food away from home" to the NAICS code "7225 Restaurants and Other Eating Places", or "Household appliances" to "3352 Household Appliance Manufacturing". However, these categories cover only goods and services (and their producing industries) that show up in the end consumers' basket. Yet for the analysis in this paper, it is important to also classify industries producing intermediate goods and services as luxuries or necessities, since these industries contribute to aggregate productivity and employment even when their output is not purchased directly by households. To extend the classification to upstream industries, I use the BEA's input-output tables and calculate an output-weighted average of the Engel elasticities of the consumption categories each industry ultimately supplies. I classify industries with sufficient coverage (at least 25% of consumption-related output) as luxuries if this average exceeds one and as necessities if it is below one.

Crossing luxury/necessity status from these CEX-estimated elasticities with goods/services status based on NAICS codes then delivers four sectors: Necessity goods, necessity services, luxury goods and luxury services. The analysis concerns reallocation between industry groups; it does not separately identify shifts toward higher-quality products within an industry. Appendix A describes the full construction and mapping details.

I classify primary-residence housing as a necessity, consistent with the combined-housing Engel estimates under both specifications (Appendix A.1). Healthcare is a necessity under Specification I (0.90) and a luxury under Specification II (1.03); it is the only category besides pet purchases for which the two specifications disagree on its luxury–necessity status. I classify it as a luxury in the baseline, motivated by Hall and Jones (2007), whose model implies that health spending rises more than proportionally with income as richer households devote a larger share of their resources to extending life.

As a robustness check, I reproduce the main plots under a classification that excludes housing and healthcare from the classified consumption categories but keeps the same Engel specification and mapping procedure for the other categories. The contrasting price, productivity, and employment paths of the two luxury sectors survive this comparison. Appendix B.1 reports these comparisons, as well as the corresponding plots under Specification II (as opposed to the Specification I baseline).

Alongside the CEX, my analysis draws on several other industry- and household-level datasets. I use the Census Bureau's Business Dynamics Statistics (BDS) to measure national employment at the four-digit industry level over 1978–2023. For value added, factor compensation, inputs, and productivity, I use the BEA–BLS Integrated Industry-Level Production Account (KLEMS), which covers 1997–2024. I merge these with historical KLEMS data starting in 1980 for the longer growth accounting sample, and I use the BEA Industry Economic Accounts to compare nominal value added between sectors. KLEMS reports data for 63 industry groups, which are coarser than the four-digit NAICS industries in my four-sector classification. When a KLEMS industry contains NAICS4 industries that are assigned to different sectors, I split its activity between the NAICS4 sub-industries using the estimated nominal-value-added shares of

its NAICS4 industries (fractional assignment). If some shares of a KLEMS industry remain unclassified, I keep their share of the outcome as unclassified rather than distributing it across the four sectors. Productivity itself is measured at the broader KLEMS industry level. Appendix A explains how I construct these weights and aggregate the data.

My baseline four-sector classification covers 297 of the 303 nongovernment four-digit NAICS2017 industries. Together, they account for 99.1% of BDS employment and 98.7% of private KLEMS nominal value added in 2023 (Appendix Table 9). For example, restaurants and hospitals are classified as luxuries, grocery stores as necessities, and employment services as an upstream luxury industry (Appendix Table 11).

I construct consumer-price indexes from D-CPI data (Jaravel 2026), using CEX expenditure shares and direct classifications of final-consumption items. The indexes measure the consumption prices of the goods and services households buy, so they do not directly cover upstream industries. I use the same four sector labels as in the employment and productivity data, but the activities covered by these labels differ as they only include final consumption here. They also measure consumer prices, while KLEMS deflators measure value-added prices. Appendix A explains how I construct the indexes and how final household expenditure differs from the domestic value added embodied in it and from total industry value added.

To connect these industry patterns to workers, I use the Current Population Survey (CPS). I take annual earnings from the Annual Social and Economic Supplement (ASEC), employment and hours from the Basic Monthly files, and hourly pay from the Outgoing Rotation Groups (ORG). I define college workers as those with a bachelor's degree or higher and include workers with associate degrees in the non-college group. Appendix A.5 describes the mapping, weights, and wage adjustments.

2.2 The two technologies of luxury

Luxury goods and luxury services have very different price, productivity and employment paths. In Figure 1, Panel A, consumer prices for luxury services more than double between 1999 and 2024, while those for luxury goods fall. Prices in both necessity sectors also increase, but less than in luxury services. Hence, over this period luxury services become more expensive relative to the other three sectors, while luxury goods become relatively cheaper.

These price patterns are accompanied by large differences in labor productivity growth (Panel B). Real value added per hour more than doubles in luxury goods between 1997 and 2024, whereas growth in luxury services is much weaker. Necessity services also have faster productivity growth than luxury services, while necessity goods have broadly similar growth to luxury services. These price and productivity patterns are consistent with technological progress reducing the relative cost of producing luxury goods.

The contrast between luxury goods and luxury services also appears with alternative price and productivity measures (Appendix Figure 20). Luxury goods have slower price growth and faster productivity growth than luxury services across these measures, although their value-added prices rise over 1999–2024.

Domestic employment shifts strongly toward luxury services (Panel C). Over 1978–2023, employment in these services almost triples, while employment in luxury goods falls by about two-fifths. Employment also expands in necessity services, but much less than in luxury services. Necessity-goods employment stays nearly flat over time. These employment patterns are consistent with differences in labor requirements across sectors: For a given increase in output, slower labor productivity growth requires a larger increase in hours worked. Growing demand for luxuries can therefore have different employment implications depending on which luxury sector supplies it. Luxury services also expand faster than the other sectors in nominal value added and firm–industry counts (Appendix Figure 21).

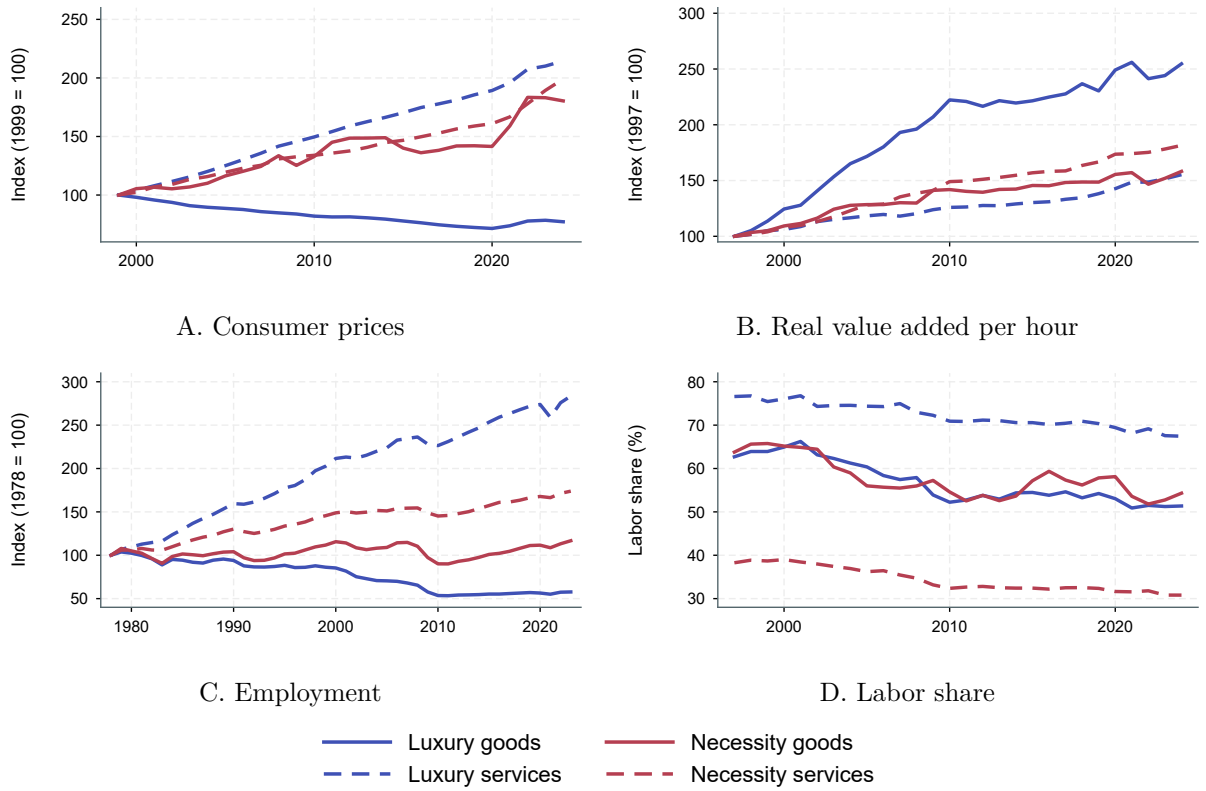
These differences also matter for labor income. Panel D shows that luxury services have the highest labor share throughout 1997–2024. Although labor shares decline over the sample, in line with the broader decline studied by Hubmer (2023), substantial differences across sectors persist. In 2024, labor compensation accounts for 67.4% of value added in luxury services, compared with around one-half in the two goods sectors and one-third in necessity services. Holding factor shares fixed, an additional dollar of value added in luxury services therefore leads to a relatively larger increase in the sector's wage bill. Whether this increases the wages of an individual worker depends on equilibrium effects and the worker's exposure to the sector's expansion. I examine which workers are exposed to this expansion in Section 2.3.

Luxury services retain the highest labor share when housing and healthcare are excluded from the classified consumption categories, although the differences narrow. In 2024, labor compensation accounts for 63.6% of their value added, compared with 52.1–54.6% in the other sectors (Appendix B.1).

The contrasting paths of the two luxury sectors could also be connected, which I explore in Section 5.4. Intuitively, productivity growth in luxury goods can raise household purchasing power by making luxury goods more affordable, and income-elastic demand can then shift a growing share of spending into luxury services. The expenditure share of luxury services can also rise because they become more expensive, even if households buy less of them. Holding the household's budget and the other sectors' prices fixed, this happens when quantities fall less than proportionally to the price increase. Hence, rising expenditure shares can reflect both income and relative-price effects. I quantify these demand responses in Section 3 to explain their contributions to the observed changes in sectoral expenditure shares. The general equilibrium model in Section 4 then connects these demand responses to sectoral employment.

One might be tempted to conclude that the overall shift from goods to services fully accounts for the reallocation toward luxuries, since services are more luxury-intensive than goods. However, this is not the case. Appendix Figure 23 holds the initial luxury shares within goods and services fixed while allowing the service share of classified activity to change as observed. This comparison path understates the observed increase in luxuries' shares of employment and nominal value added. Hence, the shift to services accounts for only part of the observed reallocation. Changes *within* goods and services matter as well: The luxury share rises within services, although it falls within goods. Appendix B.4 explains the comparison paths and their interaction.

Figure 1: Prices, productivity, employment, and labor shares



Notes: D-CPI/CEX (prices, 1999–2024), KLEMS (productivity and labor shares, 1997–2024), and BDS (employment, 1978–2023). Productivity is real value added per hour; labor shares are labor compensation divided by value added. Panels A–C are indexed to their first year.

2.3 The consumers and workers of luxury services

The high labor share of luxury services, and corresponding low capital share (Appendix Figure 24), links their expansion to labor earnings. At given factor shares, an increase in their value added generates a comparatively large increase in the sector's wage bill.

Figure 2 compares consumption and worker exposure across all four sectors. Panel A ranks households by after-tax income: Higher-income households devote larger expenditure shares to luxury services. The luxury-service expenditure share also rises from the 2000s to the 2010s within both age groups and all five education groups in Appendix Figure 28, although these comparisons do not hold all household characteristics fixed.

Panel B ranks individual workers by annual earnings. Luxury services account for a larger share of hours worked by lower-earning workers.

Luxury-service exposure declines with hourly pay within both college and non-college groups. At a given hourly-wage rank, college workers are more exposed to luxury services, and their share of hours rises toward the top of the wage distribution. The rising college share across hourly-wage ranks counteracts the declining exposure within each education group (Appendix Figure 27).

Figure 2: The consumers and workers of luxury services



A. Household expenditure by income decile

B. Worker exposure by earnings decile

Notes: Panel A: CEX nominal expenditure shares by annual after-tax household-income decile, 1980–2023. Panel B: CPS ASEC hours shares by individual annual wage-income decile, 1980–2024. Both condition on classified activity.

To connect worker exposure to the quantitative model, I rank occupations using pooled hourly pay in 1995–1997 and form five groups across all workers whose definitions remain fixed over time. Figure 3 shows that luxury services account for a larger share of classified hours in both the lowest- and highest-paid occupation groups than in the three middle groups. This pattern appears in both 1997 and 2024. Comin, Danieli, et al. (2026) document a related pattern across sectors: Those with higher expenditure elasticities use low- and high-skill occupations more intensively. An increase in the expenditure share of luxury services can therefore raise demand for both low- and high-paid occupation types. Its distributional effects across households, however, also depend on households' capital income and on what consumption baskets they buy.

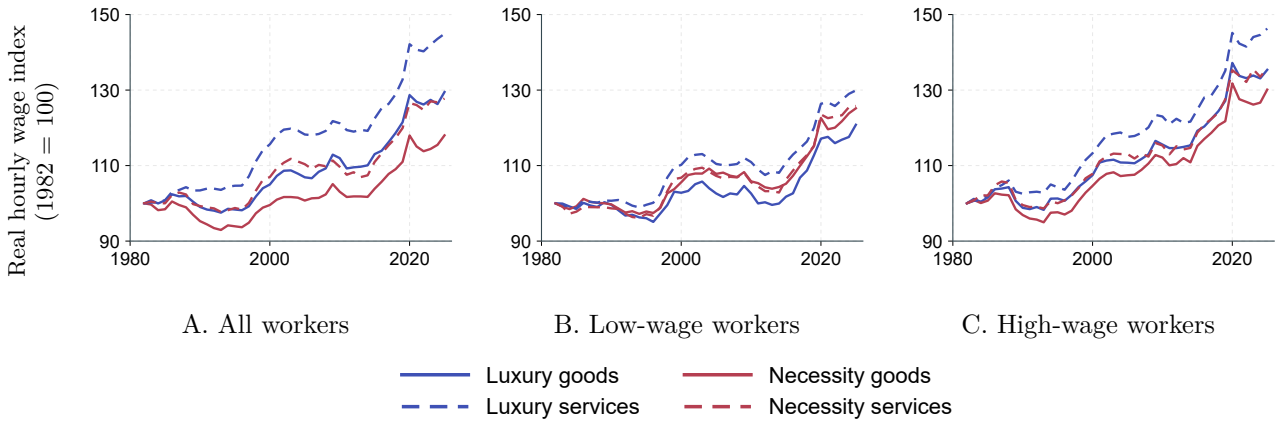
Figure 3: Luxury-service exposure by initial occupation pay



Notes: CPS Basic Monthly, 1997 and 2024; private employed workers aged 16 and above. Survey-weighted actual main-job hours, averaged over all twelve months. Occupations are ranked by pooled real hourly pay in the 1995–1997 CPS ORG and grouped using 1997 ORG hours. The five groups apply across all workers and remain fixed over time. Each bar is luxury-service hours as a percentage of the same group's total classified hours across necessity goods, necessity services, luxury goods, and luxury services in that year. Industry exposure uses the baseline fractional classification; unclassified hours are excluded.

Across all workers, luxury services have the largest increase in composition-adjusted real hourly pay (Figure 4, Panel A). Furthermore, luxury services display the strongest wage growth in both halves of each year's national hourly-wage distribution (Panels B and C). Over 1982–2025, luxury-service wages rise by 29.9% in the lower half and 46.3% in the upper half. Necessity services have the next-largest increases, at 25.8% and 35.9%, respectively.

Figure 4: Composition-adjusted real hourly wage growth



Notes: CPS ORG, 1982–2025; private employees aged 16–64. CPI-U-deflated geometric mean hourly wages; 1982 = 100. Annual wage groups split at the national median using CPS ORG survey weights. Sector means also use fractional industry exposure. Log-wage controls: Age, age squared, sex, race, Hispanic origin, state, education, and year.

I define the low-wage and high-wage groups year-by-year as a repeated cross-section, so their composition can change over time. The figures do not track individual workers' wages over time. Even after adjusting for observed worker characteristics, these comparisons do not isolate the effect of rising demand for luxury services on workers' wages.

These patterns suggest that the luxury shift can affect households differently through earnings and consumption. Higher-income households spend larger budget shares on luxury services, which makes them more exposed to increases in this sector's relative price. At the same time, growing demand for luxury services can support the wages of the occupation groups used intensively in their production. To quantify how these income and price effects jointly affect household welfare, I develop a general equilibrium model.

2.4 The luxury shift and aggregate productivity growth

Over 1980–2024, observed labor productivity growth is 0.20 percentage points per year below a fixed-weight accounting benchmark. Luxury services account for about two-thirds of this gap.

I measure labor productivity as real value added per total hour worked, using merged historical and current KLEMS data for classified private-sector activity. I first aggregate allocated nominal value added and hours to the four sectors (necessity goods, necessity services, luxury goods, and luxury services), and construct an index of real value added for each sector.

Let $\bar{s}^V_{st}$ denote sector s's midpoint share of nominal value added, Y_{st} its index of real value added, and H_{st} its total hours worked. Observed aggregate labor productivity growth equals

the weighted growth in sectoral real value added less growth in total hours:

$$g_t^{\text{obs}} = \sum_s \bar{s}_{st}^V \Delta \log Y_{st} - \Delta \log \left(\sum_s H_{st} \right). \tag{1}$$

For the benchmark, I instead average sectoral labor productivity growth using fixed nominal value-added weights from the 1980–1981 midpoint:

$$g_t^{\text{fixed}} = \sum_s \bar{s}_{s0}^V \left(\Delta \log Y_{st} - \Delta \log H_{st} \right). \tag{2}$$

Besides holding nominal value-added weights fixed, this benchmark also sets to zero the contribution to aggregate productivity growth from changes in each sector's share of total hours worked. Such shifts matter because sectors differ in nominal value added per hour. I use the same observed sectoral productivity paths in both series, so changes in industry composition within each sector enter both accounts. The gap thus shows how reallocation across the four sectors contributes to aggregate productivity growth, holding each sector's observed productivity path fixed. Appendix B.7 explains how I construct the two series and allocate the gap across sectors. These accounting results also have implications for real output growth. If aggregate hours follow the same path in both accounts, the 0.20-percentage-point annual gap in labor productivity growth translates one-for-one into a gap in real value-added growth.

Table 1 shows that annual log labor productivity growth averages 1.83% in the observed account and 2.03% in the benchmark. This difference accumulates over time, leaving observed productivity about 8.4% below the benchmark level by 2024. Figure 5 shows that luxury services contribute 0.13 percentage points per year, or about 64% of the aggregate gap. Luxury goods contribute 0.06 and necessity services 0.03, while necessity goods partly offset the gap with a contribution of −0.01 percentage points per year.

Most of the luxury-service contribution comes from hours reallocation, which accounts for about 0.09 percentage points per year (Appendix Figure 29). Intuitively, shifting a larger share of total hours toward sectors with below-average nominal value added per hour contributes to lower aggregate productivity growth relative to the benchmark.

Changes in nominal value-added weights account for a further 0.03 percentage points per year of the luxury-service contribution. Since productivity grows at different rates across sectors, changes in their weights affect aggregate growth. A shift in weight toward sectors with slower productivity growth (as is the case for luxury services), or away from those with faster growth, thus lowers aggregate growth relative to the fixed-composition benchmark. This is the main channel for luxury goods: Their declining share of nominal value added reduces the weight on a sector with relatively strong productivity growth.

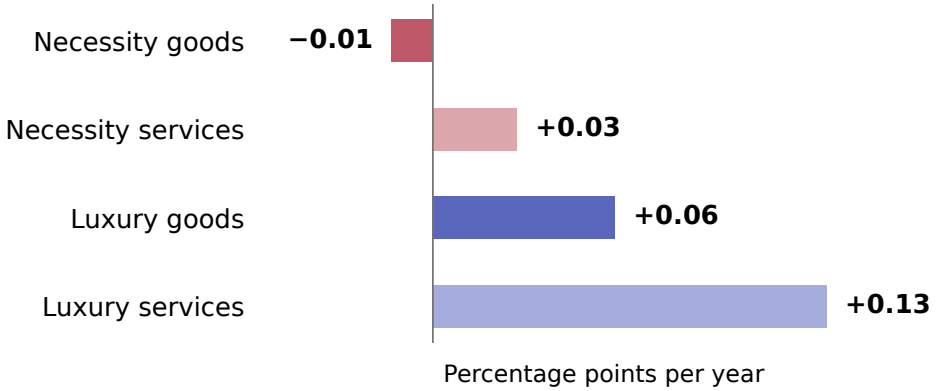
As a robustness check, I consider an industry-level comparison. This comparison also captures reallocation among the 60 KLEMS industries within each sector, whereas the baseline only compares reallocation across the four broad sectors. However, fixing value-added weights across the 60 KLEMS industries yields a very similar growth gap of 0.22 percentage points per year

Table 1: Annual labor productivity growth: Observed and fixed-weight accounts

	Annual log growth
Observed account	1.83%
Fixed-weight benchmark	2.03%
Fixed minus observed	0.20 pp

Notes: KLEMS, 1980–2024; classified private industries. Mean annual log growth of real value added per hour, in percent; the gap is in percentage points. The benchmark fixes 1980–1981 midpoint value-added weights and excludes the contribution from changes in sectors' shares of total hours worked (Appendix B.7).

Figure 5: Sectoral contributions to the labor productivity growth gap



Notes: Contributions to fixed-minus-observed growth, 1980–2024, in percentage points per year. Positive values widen the gap. Sample and attribution: Table 1 and Appendix B.7.

and essentially the same luxury-service contribution of 0.13. Appendix Table 14 compares the two accounts and documents the differences in their sectoral contributions.

The four-sector gap narrows from 0.27 percentage points per year in 1981–2004 to 0.11 in 2005–2024 (Appendix Table 13). The luxury-service contribution also falls, from 0.17 to 0.07 percentage points per year. Both periods use the original 1980–1981 midpoint weights. Hence, the productivity drag is smaller in the later period and does not account for the slowdown in productivity growth between them. Luxury services also contribute most when the four-sector account starts in 1997 or uses quality-adjusted labor (Appendix Table 15).

This accounting quantifies the contribution of changing sectoral composition to productivity growth. It does not explain *why* expenditure shifts across sectors, nor does it determine the associated household gains. To answer these questions, I first examine the demand responses behind these changes, taking household budgets and prices as given, and then allow both to adjust in general equilibrium.

3 Demand estimation and historical reallocation

In this section, I present the demand system, estimate its parameters using household data, and then use it to account for historical changes in expenditure shares. I find that income growth and relative-price changes both contribute substantially to the expansion of luxury services.

This analysis takes observed household budgets and prices as given. The demand system will be embedded in the general equilibrium model in Section 4.

3.1 The demand system

Given prices and an expenditure budget, each household chooses consumption across necessity goods, necessity services, luxury goods, and luxury services to maximize utility. I represent its nonhomothetic CES preferences, following Comin, Lashkari, et al. (2021), through the expenditure function: The resources needed to attain utility $u > 0$ at prices $\mathbf{p}$. Necessities and luxuries form the outer nest, with goods and services within each group. The nesting allows substitution between goods and services to differ within necessities and luxuries, while separately governing substitution between the two expenditure groups.

Let $j \in N, L$ index necessities and luxuries and $k \in G, S$ goods and services, so $s = jk \in \mathrm{NG}, \mathrm{NS}, \mathrm{LG}, \mathrm{LS}$ indexes the four sectors. Suppressing time subscripts, the inner expenditure indices and total expenditure function are

$$P_{ij}(\mathbf{p}_j, u) = \left[\sum_k \beta_{ijk} u^{\varepsilon_{jk}-1} p_{jk}^{1-\sigma_j}\right]^{1/(1-\sigma_j)}, \tag{3}$$

$$\mathcal{E}_i(\mathbf{p}, u) = u\left[\sum_j \beta_{ij} u^{\varepsilon_j - 1} P_{ij}(\mathbf{p}_j, u)^{1-\sigma_0}\right]^{1/(1-\sigma_0)}. \tag{4}$$

Here $\mathbf{p}_j$ contains group j's two prices. The substitution elasticities are σ_0 across necessities and luxuries and σ_N, σ_L within them. Positive taste weights β govern expenditure levels, while the ε parameters govern nonhomotheticity. I normalize $\varepsilon_N = \varepsilon_{\mathrm{NG}} = \varepsilon_{\mathrm{LG}} = 1$. The ε's are not Engel elasticities; these depend on the ε's, the substitution elasticities, and the household's current expenditure shares. So an ε above one does not by itself imply an Engel elasticity above one (i.e., that the sector is a luxury). Appendix C.1 derives this relationship.

I allow taste weights to vary with household demographics and region, and include calendar controls in the estimation. In the main historical accounting exercise, I keep the demographic and regional coefficients fixed and hold the calendar effects at their reference values. Changes in the distribution of household characteristics can thus affect aggregate demand, but changing calendar effects (e.g., changes in taste over time) cannot. The nested structure imposes three common elasticities of compensated substitution, but expenditure shares and responses to income changes can differ across households and over time.

The expenditure budget x_{it} determines attained utility, and expenditure derivatives determine optimal quantities c_{ist} and budget shares s_{ist}:

$$\mathcal{E}_i(\mathbf{p}_t, u_{it}) = x_{it}, \qquad c_{ist} = \frac{\partial \mathcal{E}_i(\mathbf{p}_t, u_{it})}{\partial p_{st}}, \qquad s_{ist} = \frac{p_{st} c_{ist}}{x_{it}}. \tag{5}$$

Income and relative prices. Nonhomotheticity changes where households spend additional resources. Within luxuries, equations (3)–(4) imply the expenditure ratio, suppressing time

subscripts,

$$\frac{s_{i,\mathrm{LS}}}{s_{i,\mathrm{LG}}} = \frac{\beta_{i,\mathrm{LS}}}{\beta_{i,\mathrm{LG}}} u_i^{\varepsilon_{\mathrm{LS}}-\varepsilon_{\mathrm{LG}}} \left(\frac{p_{\mathrm{LS}}}{p_{\mathrm{LG}}}\right)^{1-\sigma_L}. \tag{6}$$

At fixed prices, greater resources raise utility and increase luxury-service expenditure relative to luxury goods when $\varepsilon_{\mathrm{LS}} > \varepsilon_{\mathrm{LG}}$. Holding utility fixed, a higher price of luxury services relative to luxury goods raises this expenditure ratio when $\sigma_L < 1$. This is the compensated price effect. If nominal expenditure is held fixed instead, the price change also affects utility. This ratio shows how households split their luxury spending between luxury goods and luxury services. The luxury-service share out of total household expenditure also depends on how much they spend on total luxuries relative to total necessities, which is determined by the outer nest.

3.2 Estimation and Engel-curve fit

The estimation sample comes from the CEX Interview Survey for 2000–2023 and includes households with heads aged 25 or older. I require observations from all four expenditure interviews before trimming expenditure. I use the 2012 input-output tables to measure the domestic value added embodied in household purchases and allocate it across the four sectors. I use the same input-output relationships throughout the sample. To measure the value added produced domestically, I exclude imports both from households' purchases and from the intermediate inputs used by domestic firms. For consumer prices, I use regional CPI microdata to construct separate price measures for each household.

I estimate nonhomotheticity by comparing how households with different expenditure levels divide their budgets across the four sectors. I account for differences in the prices they face, household demographics, region and calendar effects. The three equations compare services with goods within necessities and luxuries, and luxuries with necessities. I instrument the expenditure terms in each equation with income-class indicators and log adjusted real after-tax household income. I estimate all demand parameters jointly by survey-weighted nonlinear GMM.

To estimate the substitution elasticities, I use relative-price variation across regions and over time. For the relative-price instruments NS/NG, LS/LG, and LG/NG, I use component CPIs from other regions, which I aggregate using contemporaneous expenditure weights from the household's own region. Identification of the demand parameters relies on the assumption that the income and price instruments are uncorrelated with unobserved taste shocks, conditional on the controls. I explain the estimating equations and how I construct the instruments in Appendix C.2.

In the baseline, I trim real sectoral expenditure and original total expenditure at the 2nd and 98th percentiles within each year. The baseline specification, reported in column (3) of Table 2, includes household demographic controls, region fixed effects, and year-by-quarter indicators, following Comin, Lashkari, et al. (2021). Column (2) omits calendar controls, and column (1) also omits region effects. All three specifications imply substitution elasticities below one at every nest. The fit of the Engel curves is also robust to these changes. Specification (3) uses 269,518 interviews from 71,938 households, with household-clustered standard errors.

Table 2: Estimated demand parameters: Domestic consumption value added

Parameter / specification	(1)	(2)	(3)
σ_0	0.44	0.46	0.41
	(0.01)	(0.01)	(0.02)
σ_N	0.17	0.31	0.20
	(0.02)	(0.02)	(0.03)
σ_L	0.21	0.19	0.24
	(0.01)	(0.01)	(0.01)
ε_L	1.42	1.56	1.41
	(0.02)	(0.03)	(0.03)
$\varepsilon_{\mathrm{NS}}$	1.28	1.36	1.29
	(0.01)	(0.02)	(0.02)
$\varepsilon_{\mathrm{LS}}$	1.81	1.94	1.81
	(0.02)	(0.03)	(0.03)
Region fixed effects	No	Yes	Yes
Region-specific linear trends	No	No	No
Calendar controls	No	No	Yes
Estimation weight	Survey	Survey	Survey
Observations	269,518	269,518	269,518
Households	71,938	71,938	71,938

Notes: CEX, 2000–2023; survey-weighted nonlinear GMM for domestic consumption value added, using regional price and household income instruments. Column (3) is preferred and includes calendar controls. Household-clustered standard errors in parentheses. Normalizations: $\varepsilon_N = \varepsilon_{\mathrm{NG}} = \varepsilon_{\mathrm{LG}} = 1$.

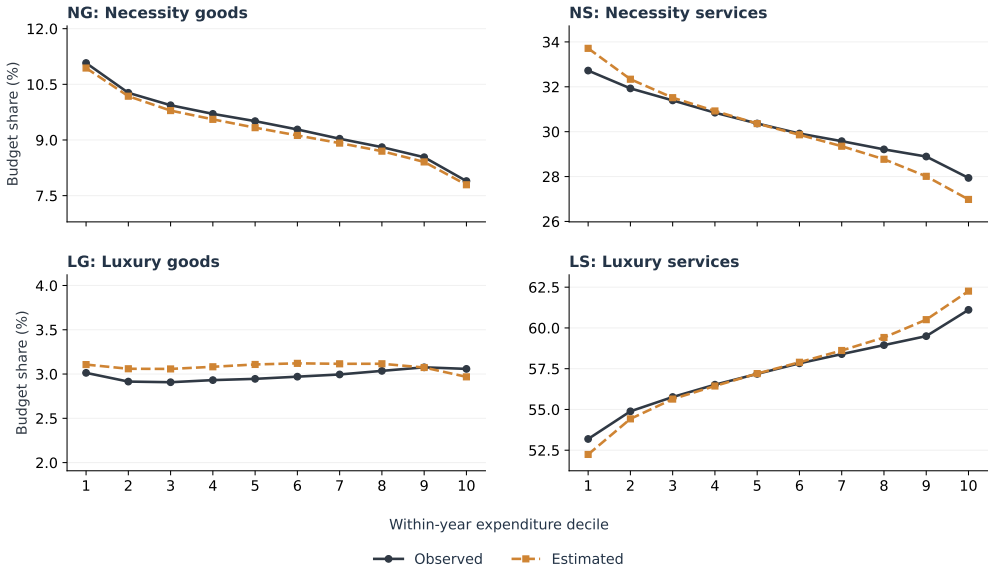
I use column (3) as the baseline because its time fixed effects allow for aggregate changes in sectoral demand over time when estimating demand responses. In the main historical demand accounting exercise, I hold these effects fixed at their omitted reference values. This means that the predicted reallocation only reflects changes in observed budgets, prices, and household composition under a fixed estimated demand relationship. Appendix C.6 allows the estimated calendar effects to change over time while keeping the same initial-share calibration. Changing calendar effects account for only 0.22 percentage points, or about 8%, of the observed increase in the luxury-service share over 1999–2023, with income and relative-price effects remaining the main contributors.

The preferred estimates imply limited substitution at all three nests: $\sigma_0 = 0.41$, $\sigma_N = 0.20$, and $\sigma_L = 0.24$. Within luxuries, a 1% increase in the relative price of services reduces their relative quantity by about 0.24% at fixed utility. Spending on luxury services relative to luxury goods therefore rises by about 0.76%, as the decline in their relative quantity is smaller than the increase in their relative price.

Figure 6 shows that the estimated system captures the main differences in spending patterns across expenditure deciles. The luxury-service share rises across deciles, necessity shares decline, and the luxury-goods share changes little. For these Engel-curve predictions, I evaluate the estimated calendar effects at each observation date. For the historical demand accounting below, however, I keep the calendar effects fixed at their omitted reference values. This way, the predicted changes in expenditure shares reflect only the estimated demand responses to changes in observed household budgets, prices and characteristics. Appendix Figure 30 shows observed and predicted shares against real-dollar expenditure. Including imports changes the estimated

substitution elasticities, but they remain below one at all three nests, and the system continues to reproduce most of the historical luxury-service expansion in the corresponding consumption measure (Appendix C.4).

Figure 6: Observed and predicted Engel curves



Notes: CEX, 2000–2023; survey-weighted domestic consumption-VA shares by within-year expenditure decile. Predictions use column (3) of Table 2, household-specific prices and tastes, and the estimated calendar effects at each observation date.

3.3 Historical demand accounting

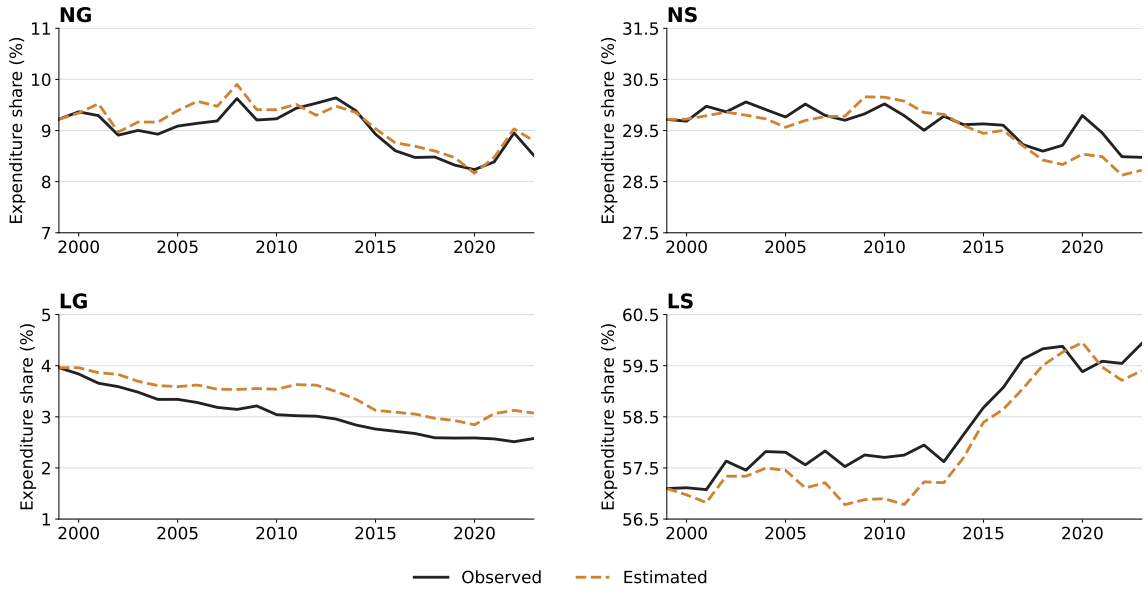
The Engel curves show how well the system captures differences in spending patterns across households. I now ask how much of the change in aggregate expenditure shares over time these estimated demand responses can account for. With calendar effects held at their omitted reference, the system predicts a 2.31-percentage-point rise in the luxury-service share over 1999–2023, compared with 2.84 points in the data (Figure 7). Both income and relative-price changes contribute substantially.

I calibrate three taste intercepts to match the aggregate shares in 1999 and keep these intercepts and all estimated coefficients fixed thereafter. I then predict expenditure shares in the following years only from observed household budgets, prices, and demographic and regional characteristics. Thus, I match the initial shares by construction and use the estimated demand relationship to predict changes in the years after that.

The accounting sample applies the matched within-year 2nd–98th percentile trims to raw expenditure. Annual sectoral prices are survey-weighted geometric means of household price exposures for domestic consumption value added (Appendix C.5).

The system also captures the decline in goods shares, although it understates the luxury-goods decline and overstates the necessity-service decline.

Figure 7: Observed and predicted sectoral expenditure shares



Notes: CEX, 1999–2023; domestic consumption value added. Predictions use observed budgets, prices, and composition, with three taste intercepts calibrated to 1999 shares and fixed in the years after that. Estimated coefficients are held fixed and calendar effects are kept at their omitted reference. Both series (observed and predicted) use the within-year 2nd–98th percentile sample.

To obtain aggregate expenditure shares, I weight each household's predicted budget shares by its share of total spending. Let s_{is} be household i's predicted budget share in sector s, x_i its expenditure budget, and π_i its population weight, normalized to sum to one. The predicted aggregate share and household expenditure weights are therefore

$$\widehat{S}_s = \sum_i \nu_i s_{is}, \qquad \nu_i = \frac{\pi_i x_i}{\sum_h \pi_h x_h}. \tag{7}$$

To decompose changes in these aggregate shares over time, I link adjacent annual cross-sections by population expenditure ranks, preserving each year's joint distribution of budgets and characteristics. Each rank interval represents a fixed share of the population along its link. I compare households at the same expenditure rank across years, rather than follow the same households over time.

On this common population-rank representation, the aggregate share change has four components:

$$\begin{aligned} d\widehat{S}_s = {} & \underbrace{\mathbb{E}_\nu[s_{is} B_{is}]}_{\text{Compensated relative-price effects}} \\ & + \underbrace{\mathbb{E}_\nu\left[s_{is}(\eta_{is}-1)\left(d\log x_i - \sum_r s_{ir}\, d\log p_r\right)\right]}_{\text{Income effects}} \\ & + \underbrace{\mathrm{Cov}_\nu(s_{is}, d\log x_i)}_{\text{Expenditure reweighting}} + \underbrace{\mathbb{E}_\nu[T_{is}]}_{\substack{\text{Demographic and regional} \\ \text{composition}}} . \end{aligned} \tag{8}$$

Expectations and covariances use expenditure weights ν_i. Here $\eta_{is} = \partial \log c_{is} / \partial \log x_i$ is the household's expenditure elasticity at fixed prices and characteristics. The compensated response B_{is} measures how log expenditure shares change with prices, holding utility and household characteristics fixed. Income effects capture how households change what sectors they spend on as their budgets grow, after adjusting for inflation in their consumption baskets. Expenditure reweighting reflects changes in each household's share of total spending and therefore in how much each household matters for aggregate demand. Finally, T_{is} measures how changes in demographic and regional composition affect the level of the household's expenditure share. I calculate this term holding prices, the household's budget, and the estimated coefficients fixed.

The nested preferences make the compensated price response explicit. For sector s in group $j \in \{N, L\}$,

$$B_{is} = (1 - \sigma_0)\left(d \log \mathcal{P}_{ij} - d \log \mathcal{P}_i\right) + (1 - \sigma_j)\left(d \log p_s - d \log \mathcal{P}_{ij}\right), \tag{9}$$

where

$$d \log \mathcal{P}_i = \sum_r s_{ir}\, d \log p_r, \qquad d \log \mathcal{P}_{ij} = \sum_{r \in j} \frac{s_{ir}}{\sum_{m \in j} s_{im}}\, d \log p_r. \tag{10}$$

These are the overall and within-group expenditure-weighted price-change components. They hold utility and household characteristics fixed, so they differ from total changes in the utility-dependent expenditure indices in equation (3). The first term in equation (9) captures substitution between necessities and luxuries; the second captures substitution between goods and services within each group.

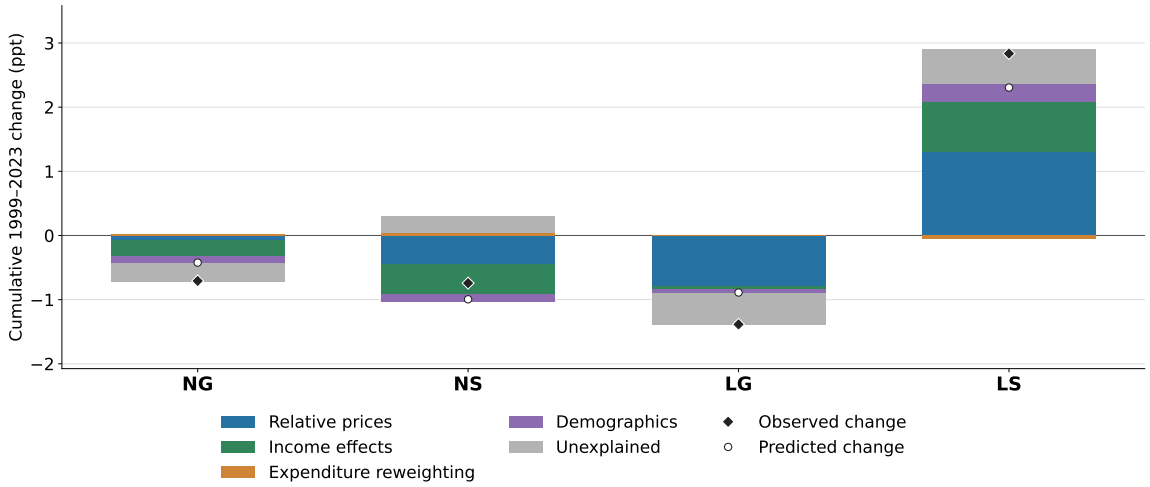
I construct the contribution bars by integrating the four terms between consecutive years and summing their contributions over 1999–2023. Let $\mathcal{C}_s^P$, $\mathcal{C}_s^I$, $\mathcal{C}_s^R$, and $\mathcal{C}_s^D$ denote the resulting percentage-point contributions from relative prices, income, expenditure reweighting, and composition. With $\mathcal{U}_s = 100(\Delta S_s^{\text{observed}} - \Delta \widehat{S}_s)$, the complete reconciliation is

$$100 \Delta S_s^{\text{observed}} = \underbrace{\mathcal{C}_s^P + \mathcal{C}_s^I + \mathcal{C}_s^R + \mathcal{C}_s^D}_{100 \Delta \widehat{S}_s} + \underbrace{\mathcal{U}_s}_{\text{Unexplained change}}. \tag{11}$$

The components depend on the specified accounting path. These contributions are not separately identified exogenous shocks. They differ from the general-equilibrium experiments in Section 5, where I change preferences or technology and solve for the resulting household budgets and prices. Appendix C.5 derives the identity and explains the integration.

Relative prices and income are both quantitatively important (Figure 8). Compensated price effects contribute 1.30 percentage points to the luxury-service expansion, and income effects add 0.78 points. Limited substitution allows their expenditure share to rise as they become relatively more expensive; nonhomotheticity raises it further as purchasing power grows. Composition adds 0.28 points and expenditure reweighting offsets 0.06 points, completing the predicted 2.31-point increase. A further 0.53 points of the observed increase remain unexplained by this account. Appendix Figure 34 shows the contributions over time.

Figure 8: Sources of sectoral expenditure reallocation



Notes: Changes in domestic consumption-value-added shares, 1999–2023, in percentage points. Prices are compensated effects; income is net of household-basket inflation. Expenditure reweighting captures changes in households' shares of total spending; demographics also includes regional composition. Unexplained is the observed change minus the predicted change. I hold the estimated taste coefficients fixed and calendar effects at their omitted reference.

Average growth and expenditure inequality. I ask how much of the income contribution to luxury services reflects higher average spending and how much reflects changes in expenditure inequality. Along the same historical path, I split each household's expenditure growth into growth in the mean and changes in its expenditure relative to that mean. The mean component includes the adjustment for inflation in each household's consumption basket and accounts for 0.77 of the 0.78-percentage-point income contribution. Changes in relative expenditure add only 0.01 points through Engel responses (Appendix Figure 35). Changes in the expenditure distribution also change how much weight each household has in aggregate spending. I account for this separately as expenditure reweighting, and it reduces the luxury-service share by 0.06 percentage points. Hence, the mean expenditure component accounts for almost all of the income contribution, while changes in expenditure inequality contribute very little to the income effect on luxury-service expansion.

I also ask how expenditure shares would change if average real expenditure and its distribution changed while sector prices and household characteristics stayed at their 1999 values (Appendix Figure 36). Here I deflate expenditure using a common CPI and allow both household budget shares and aggregate spending weights to adjust. Changing mean expenditure alone produces almost the same rise in the luxury-service share as changing both the mean and the distribution. Changing only the distribution has little effect on that share.

These contributions show how observed budgets and prices translate into changes in domestic consumption-value-added shares. The general equilibrium model in Section 4 compares economies with different preferences and productivity paths to quantify how these forces affect reallocation and household welfare through changes in incomes and prices.

4 A general equilibrium model of the luxury shift

I use a four-sector general equilibrium model to quantify how nonhomothetic demand and uneven productivity growth contribute to the expansion of luxury services and affect household welfare. The preferences estimated in Section 3 determine how households change their expenditure shares as their resources grow. Sectoral differences in productivity growth and factor intensities affect relative prices and the demand for capital and different types of labor. Households differ in both how much capital they own and the occupation types of labor they supply. Hence, the welfare effects of the luxury shift depend on both how households earn their income and what they consume.

The economy has four competitive sectors, $s \in \{\text{NG}, \text{NS}, \text{LG}, \text{LS}\}$: Necessity goods, necessity services, luxury goods, and luxury services. Each produces a value-added composite. A sequence of static equilibria follows exogenous historical paths for technologies and aggregate factor supplies. Household population weights and shares of each factor endowment remain fixed; aggregate capital and labor quantities can change, along paths common across counterfactuals. Capital and worker types move across sectors, with goods prices and factor returns clearing markets. The model isolates intratemporal adjustment, abstracting from saving choices, endogenous capital accumulation, occupational switching, and endogenous labor supply.

4.1 Household resources and financing

Household i has population weight π_i, with $\sum_i \pi_i = 1$. It owns capital $\bar{K}_{it}$ and supplies hours $\bar{n}_{ijt}$ of worker type $j \in \{1, \ldots, 5\}$. Types are fixed occupation groups ordered by initial occupation pay. A household can supply more than one type of labor because its members can work in different occupations. At rental rate r_t and hourly type wages w_{jt}, gross factor income and endowment paths satisfy

$$y_{it}^F = r_t \bar{K}_{it} + \sum_j w_{jt} \bar{n}_{ijt}, \qquad \bar{F}_{it} = \bar{F}_{i0} \frac{\bar{F}_t}{\bar{F}_0}, \quad F \in \{K, n_j\}. \tag{12}$$

Fixed ownership shares therefore do not imply fixed income shares: Factor returns adjust in equilibrium.

The gap between initial household expenditure and factor income in the data can reflect transfers, saving and dissaving, omitted resources, and measurement differences (all of which are not part of the model). I reconcile this gap through a balanced financing rule. Each payer contributes a fixed fraction τ_i of factor income, and each recipient receives a fixed share of aggregate payments through coefficient b_i. Disposable resources are

$$y_{it} = (1 - \tau_i) y_{it}^F + b_i \mathcal{R}_t, \qquad \mathcal{R}_t = \sum_i \pi_i \tau_i y_{it}^F, \qquad \sum_i \pi_i b_i = 1. \tag{13}$$

I calibrate these coefficients initially and hold them fixed thereafter. The amounts households pay and receive change with their factor incomes, and the total financing payments equal the total receipts in every period. How much the recipients receive therefore depends on the (factor)

income of the contributors. Hence, changes in wages and capital returns affect recipients both directly through their own factor incomes, as well as indirectly through the payments they receive. Appendix D.2 explains how I construct the rule and scale all households' initial expenditure by the same factor.

Household demand follows equations (3)–(5), with the expenditure budget set equal to disposable resources, $x_{it} = y_{it}$. In the quantitative equilibrium model, common taste weights match initial producer-value-added shares and remain fixed. Households still consume different baskets because their budgets differ and demand is nonhomothetic.

4.2 Production and labor demand

Sectors differ in productivity, capital intensity, and the mix of occupations they employ. I distinguish five occupation groups so that changes in sectoral demand can affect their wages differently. Workers can move freely across the four sectors, but they cannot change their type. This means that each type earns the same market-clearing wage in every sector. Sectoral average pay differs then only because sectors employ different combinations of types.

Sector s produces output Y_{st} from capital K_{st} and a single labor composite ℓ_{st}:

$$Y_{st} = A_{st} K_{st}^{\alpha_{st}} \ell_{st}^{1-\alpha_{st}}, \tag{14}$$

$$\ell_{st} = \prod_{j=1}^{5} (\psi_t n_{sjt})^{a_{sjt}}, \qquad \sum_{j=1}^{5} a_{sjt} = 1. \tag{15}$$

Here A_{st} is Hicks-neutral productivity, α_{st} the capital share, and a_{sjt} the weight on worker type j. I set the elasticity of substitution across worker types to one in the baseline, so the labor aggregate is Cobb–Douglas. Appendix D.1 contains the general CES specification. ψ_t converts hours worked into effective labor services, and it is the same across all worker types and sectors at a given date. Hourly pay is $w_{jt} = v_{jt}\psi_t$, where v_{jt} is the return per effective unit of labor. Labor quality and production weights follow calibrated historical paths described in Section 4.5.2, common across counterfactuals.

Prices and factor payments. Let W_{st} be the unit cost of the labor composite. Cost minimization gives $W_{st} = \prod_j (v_{jt}/a_{sjt})^{a_{sjt}}$. Competition and constant returns imply that prices equal unit costs:

$$p_{st} = \frac{r_t^{\alpha_{st}} W_{st}^{1-\alpha_{st}}}{A_{st} \alpha_{st}^{\alpha_{st}} (1-\alpha_{st})^{1-\alpha_{st}}}. \tag{16}$$

A sector's price therefore depends on both its productivity and the equilibrium costs of capital and labor. Write $X_{st} = p_{st} Y_{st}$ for sectoral revenue. Factor payments satisfy

$$r_t K_{st} = \alpha_{st} X_{st}, \qquad w_{jt} n_{sjt} = (1-\alpha_{st}) a_{sjt} X_{st}. \tag{17}$$

The labor share determines how much of sectoral revenue is paid to workers. With unit elasticity of substitution, each occupation type's production weight equals its share of the sector's wage

bill. These weights can vary over time, but I hold them fixed across counterfactual economies at a given date.

4.3 Competitive equilibrium

Aggregate endowments are $\bar{K}_t = \sum_i \pi_i \bar{K}_{it}$ and $\bar{n}_{jt} = \sum_i \pi_i \bar{n}_{ijt}$. Given technologies, factor supplies, preferences, and financing coefficients, an equilibrium consists of prices, factor returns, household consumption, and sectoral allocations satisfying household optimality, firm cost minimization with prices equal to unit costs, and market clearing:

$$Y_{st} = \sum_i \pi_i c_{ist}, \qquad \sum_s K_{st} = \bar{K}_t, \qquad \sum_s n_{sjt} = \bar{n}_{jt}. \tag{18}$$

The effective-labor return of the fifth type is fixed at its initial value as numeraire, $v_{5t} = v_{50}$. Within the model, market clearing equates sectoral expenditure, revenue, and value added. Aggregate expenditure and sectoral expenditure shares therefore satisfy

$$\begin{aligned} X_t &\equiv \sum_s X_{st} = \sum_i \pi_i y_{it}, \qquad \nu_{it} \equiv \frac{\pi_i y_{it}}{X_t}, \\ S_{st} &\equiv \frac{X_{st}}{X_t} = \sum_i \nu_{it} s_{ist}. \end{aligned} \tag{19}$$

The population weights π_i are fixed, whereas expenditure weights ν_{it} change with after-financing budgets y_{it}. A sector's share of aggregate spending can increase either because households spend a larger share of their budgets on that sector, or because expenditure grows faster among households that already devote a larger budget share to it.

I construct a Divisia price index P_t by weighting log sectoral price changes with expenditure shares along each economy's equilibrium path (Appendix D.3). Each economy uses its own prices and shares, with P_t initially set to 100. Aggregate real output is nominal value added X_t deflated by this price index. Since aggregate hours follow the same path across counterfactual economies, differences in real output growth between them equal differences in aggregate labor productivity growth. I report real hourly wages, the real rental rate of capital, and relative sectoral prices as w_{jt}/P_t, r_t/P_t, and p_{st}/P_t; plotted indexes are rebased to 100 in their first displayed year. This reporting index is distinct from the equilibrium numeraire. I use household expenditure functions for exact welfare comparisons because households consume different baskets.

4.4 Factor returns and household gains

Factor demand and returns. Summing equation (17) across sectors and imposing factor-market clearing gives

$$\Lambda_{jt} = \sum_s (1 - \alpha_{st}) a_{sjt} S_{st}, \tag{20}$$

$$w_{jt}\bar{n}_{jt} = X_t \Lambda_{jt}, \qquad r_t \bar{K}_t = X_t \sum_s \alpha_{st} S_{st}. \tag{21}$$

Here Λ_{jt} is the fraction of aggregate expenditure paid to worker type j. At a given date, relative type supplies are the same across counterfactuals. Differences in relative wages therefore reflect differences in the shares of aggregate expenditure paid to each type. Initial pay alone does not determine how much each type is exposed to luxury-service demand: Both low- and high-paid occupation groups can be (and indeed are) concentrated in this sector. Since each occupation type earns the same wage across sectors, changes in demand for that type also affect workers of the same type employed outside of luxury services.

Productivity growth in other sectors can release labor and raise households' purchasing power. Nonhomothetic demand then shifts part of this additional spending toward luxury services. How employment and average pay in this sector change depends on the types of labor it uses and how wages and capital returns adjust in equilibrium. Let $n_{st} = \sum_j n_{sjt}$ be sectoral hours and $\bar{w}_{st}$ the wage bill per hour. The sector's wage bill satisfies

$$\bar{w}_{st} n_{st} = (1 - \alpha_{st}) X_{st}. \tag{22}$$

For a given labor share, higher spending can increase hours worked, average pay, or both. Because aggregate labor supplies are the same across counterfactuals at a given date, any additional hours worked in luxury services must come from other sectors. Hence, even a strong demand response to rising incomes does not guarantee higher wages for every low-paid worker.

Household gains. For fixed preferences, the local change in purchasing power combines disposable-resource growth and household-specific price changes:

$$d \log y_i - \sum_s s_{is} d \log p_s. \tag{23}$$

Changes in wages and capital returns affect household resources through labor and capital ownership and the financing rule. Households' expenditure shares determine how exposed they are to price changes. A household can therefore gain from higher demand for the labor that it supplies, while at the same time spending little on the sector's output. This leaves it less exposed to rising prices in that sector. For finite changes, I use the exact comparison in equation (28) rather than this local expression. I value budget opportunities using the same reference preferences and prices and express the difference relative to each household's initial budget.

4.5 Calibration

Table 3 gives an overview of the model calibration. It separates the demand parameters I estimate, the historical inputs I take from the data, the objects I calibrate, and the outcomes that are untargeted. I use historical changes in relative wages between the five occupation types to calibrate time paths of common labor-demand shifts. In each counterfactual, I let wages adjust in equilibrium rather than recalibrating them to match the historical data. The counterfactuals therefore show how changes in preferences (homothetic versus nonhomothetic) or sectoral productivity (even versus uneven across sectors) affect wages and households' purchasing power, while keeping the same calibrated shifts in demand for each occupation type.

Table 3: Model inputs, calibration targets, and untargeted outcomes

Role	Objects
Estimated demand	Six shape parameters from the preferred household demand specification.
Observed inputs	Reference VA-TFP, capital shares and aggregate factor-service paths from KLEMS; CPS occupation hours; initial household expenditure and factor holdings.
Calibrated objects	Initial tastes match producer-VA shares; production weights match reconciled wage bills by sector and worker type; financing matches initial budgets. Four common annual labor-demand shifts match relative occupation wages.
Untargeted outcomes	Later sectoral expenditure and hours shares, hours by sector and worker type, relative goods prices, absolute real-wage growth, counterfactual allocations, and household welfare.

Notes: I set the elasticity of substitution across worker types to one. Initial sectoral hours enter the wage-bill reconciliation, so they serve as calibration checks rather than independent validation.

4.5.1 Households, worker types, and parameterization

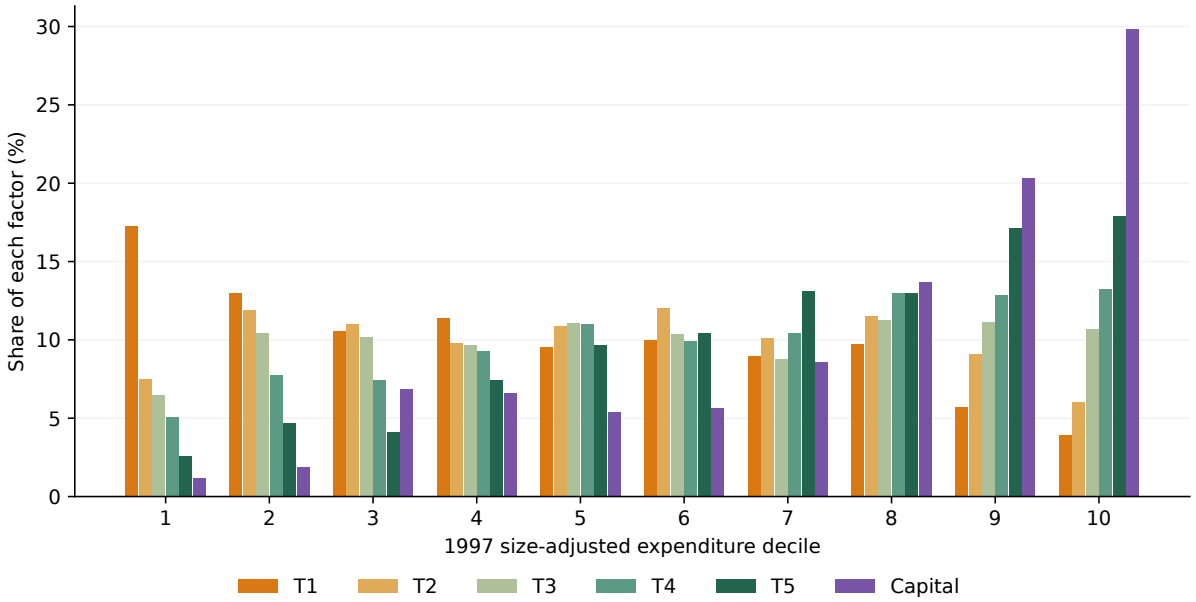
The historical calibration covers 1997–2024. I transfer the six demand shape parameters from the preferred domestic consumption-VA estimate in column (3) of Table 2 and calibrate common initial taste levels to producer-VA shares. Domestic consumption-VA shares differ from producer-VA shares because domestic industries also produce for investment and other final uses. For the model, I assume that the demand system estimated from household consumption applies to producer VA, even though the consumption-VA data do not tell me how demand for other final uses behaves. I keep the common taste levels fixed after the initial year and leave the estimated calendar effects out of the model. Subsequent sectoral shares are therefore model outcomes.

The household distribution combines 2,496 CEX households from 1997 with capital-ownership information from the 1998 Survey of Consumer Finances. I keep this capital match unchanged when constructing the occupation-specific labor holdings. To account for the occupations of different workers within a household, I use household replicas that represent their joint occupation assignments. Each replica has the original household's capital, total hours and initial expenditure. I assign weights to the 13,830 replicas so that together they preserve each household's original weight and the aggregate supply of each labor type. I keep these assignments and weights fixed across all experiments. Appendix D.2 explains the matching and financing assumptions.

Figure 9 shows ownership of each of the five types' hours and capital across the same initial expenditure deciles used for welfare. Each series sums to 100% across deciles. Thus, the figure describes where each factor is held, rather than each decile's portfolio composition. Different types of labor reach households throughout the expenditure distribution, whereas capital ownership is concentrated toward the top. A U-shaped pattern of gains across occupation types therefore need not imply a U-shaped pattern across household expenditure deciles.

I rank occupations using pooled 1995–1997 hourly pay across all workers and form five groups using the initial hours distribution. These occupation definitions remain fixed over time. I adjust the raw CPS series for the 2003 change in CPS occupation classification (from 1990

Figure 9: Initial factor ownership by size-adjusted expenditure decile



Notes: CEX 1997, CPS occupation assignments, and SCF 1998; shares of aggregate holdings before financing. T1–T5 are fixed occupation groups ordered by initial occupation pay. Household-weighted deciles rank observed 1997 consumption-VA expenditure divided by the square root of household size. Each factor's shares sum to 100%.

Census occupation classification to the 2002 Census classification). I use occupations observed both before and after the coding change to link wage growth across the break. This prevents the change in occupation codes from showing up as a shift in labor demand. The unit production elasticity is an assumption, distinct from the three estimated demand elasticities. Appendix E.7 compares five and ten groups and elasticities of one and 1.5.

The initial labor share is 77% in luxury services, compared with 64% in necessity goods, 63% in luxury goods and 38% in necessity services (Table 4). These differences mean that the implications of the overall shift toward services for the aggregate labor share depend on which services expand. In particular, a shift in spending from necessity to luxury services increases the share of income paid to labor, holding sectoral factor shares fixed. A conventional goods/services split captures the overall expansion of services, but cannot distinguish this reallocation within services, which matters for the demand-driven changes in the aggregate labor share as studied by Hubmer (2023).

4.5.2 Production paths and labor-demand shifts

KLEMS disciplines capital services, aggregate labor services, and reference sectoral productivity. CPS occupation hours determine relative type-supply growth. I scale their paths to the observed aggregate hours path per household and use one common quality adjustment to reconcile their compensation-weighted quantity index with aggregate KLEMS labor services. In particular,

$$d\log\psi_t = d\log L_t^{\text{KLEMS}} - \sum_j \omega_{jt}^L \, d\log \bar{n}_{jt}, \qquad \omega_{jt}^L = \frac{w_{jt}^{\text{target}} \bar{n}_{jt}}{\sum_h w_{ht}^{\text{target}} \bar{n}_{ht}}. \tag{24}$$

Table 4: Quantitative model: Parameters and targets

Object	Value	Source or target
$\sigma_0, \sigma_N, \sigma_L$	$0.41, 0.20, 0.24$	Household demand estimate
$\varepsilon_L, \varepsilon_{\mathrm{NS}}, \varepsilon_{\mathrm{LS}}$	$1.41, 1.29, 1.81$	Household demand estimate
Worker types J	5	Fixed pooled occupation groups
Labor substitution η	1	Cobb–Douglas baseline
Common annual shifts	4	Adjusted ORG relative wages
Initial type weights	4×5	Reconciled sector/type labor bills
Initial capital shares	$0.36, 0.62, 0.37, 0.23$	KLEMS, NG/NS/LG/LS
Households / joint replicas	2,496/13,830	CEX 1997, fixed assignments
Labor-services growth per household	2.26%	KLEMS, 1997–2024

Notes: Capital shares use KLEMS factor-payment ratios. I calibrate the initial occupation weights to match reconciled labor bills by sector and worker type. I start the reconciliation from CPS hours and ORG wages. The model matches the full matrix of reconciled labor bills and the initial producer value-added shares. I keep production and aggregate factor-supply paths the same across counterfactuals.

The adjustment counts changing occupation composition once and reproduces 2.26% growth in the reference labor-services index per household over 1997–2024. It includes measurement and coverage differences between CPS and KLEMS; it is not an identified measure of individual productivity.

Labor-demand weights evolve through common type-specific shifts:

$$a_{sjt} = \frac{a_{sj0} \exp(g_{jt})}{\sum_{h=1}^{5} a_{sh0} \exp(g_{ht})}, \qquad g_{5t} = 0, \quad g_{j0} = 0. \tag{25}$$

There are four independent annual shifts, shared across the four sectors. They match four composition-adjusted ORG relative wage-growth paths, using type 5 as reference. The initial weights differ across sectors, so common shifts can still produce different changes in their normalized weights. However, changes in type-pair log weight ratios are the same in every sector. Later sector-by-type hours are not targeted.

This restriction keeps persistent sectoral differences in labor use while allowing relative occupation demand to change over time. It allows the model to match historical relative occupation wages while keeping the cross-sector pattern of relative type demand fixed. Hence, it cannot reproduce all changes in relative type intensities across sectors. In all counterfactuals, the calibrated paths of labor-demand shifts remain unchanged, and relative wages are solved again without matching the historical targets.

The reference KLEMS VA-TFP index T_{st}^{ref} differs from Hicks-neutral productivity A_{st} because changing capital shares and type weights also contribute to measured productivity growth. I subtract their cumulative contribution along the calibrated reference factor-price path to avoid counting it twice. Realized sectoral TFP growth in the counterfactuals can nevertheless differ from the baseline because the contribution of changing production weights depends on equilibrium relative factor returns. Appendix D.3 explains the normalization.

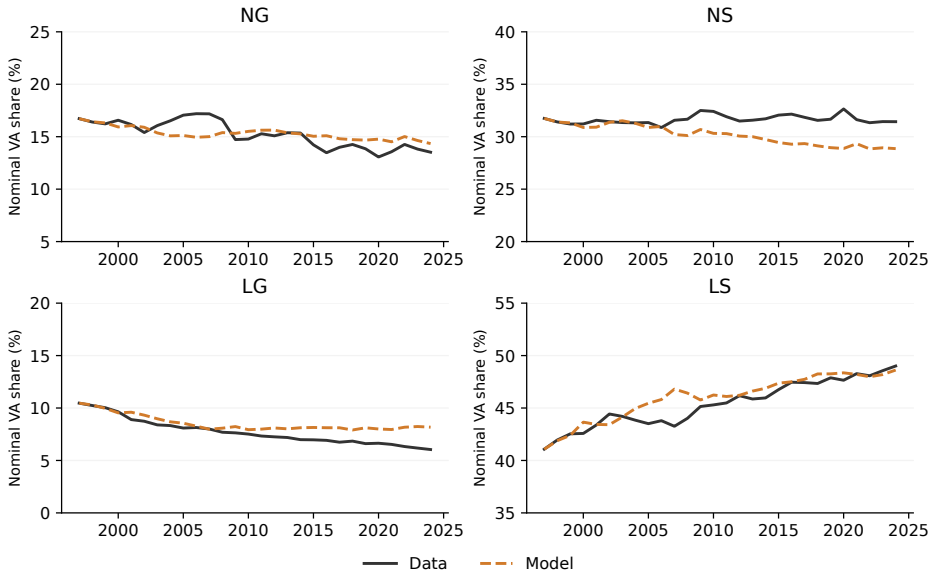
4.6 Model fit

Figure 10 shows that the model reproduces most of the increase in the luxury-service share of nominal value added. It also generates a rising share of hours in luxury services, although the increase is slightly larger than in the data. The fit is weaker for necessity services: Their value-added share remains roughly constant in the data, while the model predicts a small decline. The model also predicts a slightly larger drop in necessity-service hours than in the data, and a smaller decline in both the value-added and hours shares of luxury goods.

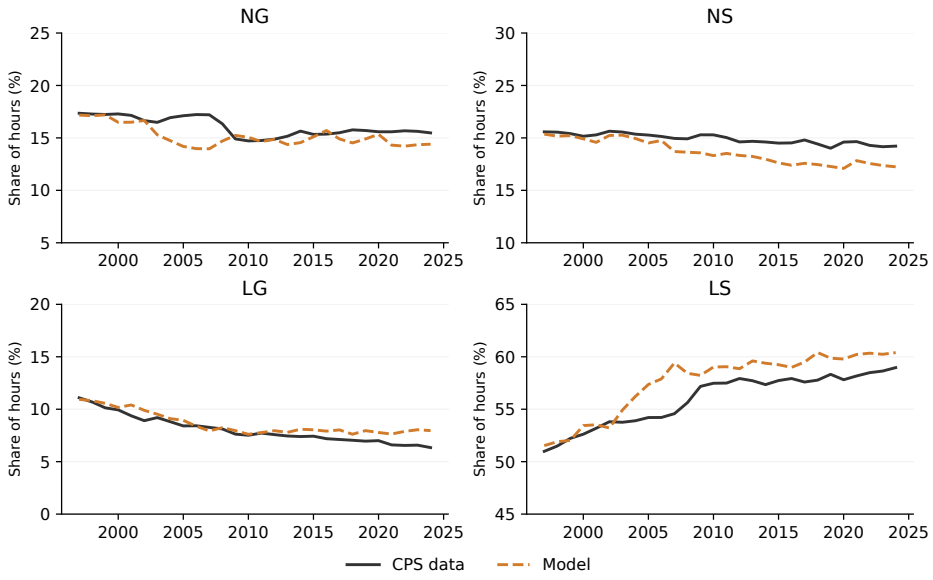
The initial value-added shares are matched by construction, and initial hours enter the labor-bill calibration. The sectoral value-added and employment shares and relative prices in later years, however, are untargeted outcomes. They are determined conditional on the production and aggregate factor-supply paths that I impose from the data.

Figure 11 compares relative producer prices in the model and the data. The model captures well the increase in the relative price of luxury services and the initial fall in that of luxury goods. However, luxury goods continue to become even cheaper after the late 2000s in the data, which the model does not capture. Necessity-service prices remain fairly stable in both, while necessity-good prices fluctuate more in the data, as this sector contains goods with high price volatility like fuels and groceries. Each set of prices is divided by its own aggregate price index. These are value-added prices, distinct from the consumer prices used in demand estimation.

Figure 10: Sectoral reallocation in the data and model



A. Nominal value-added shares

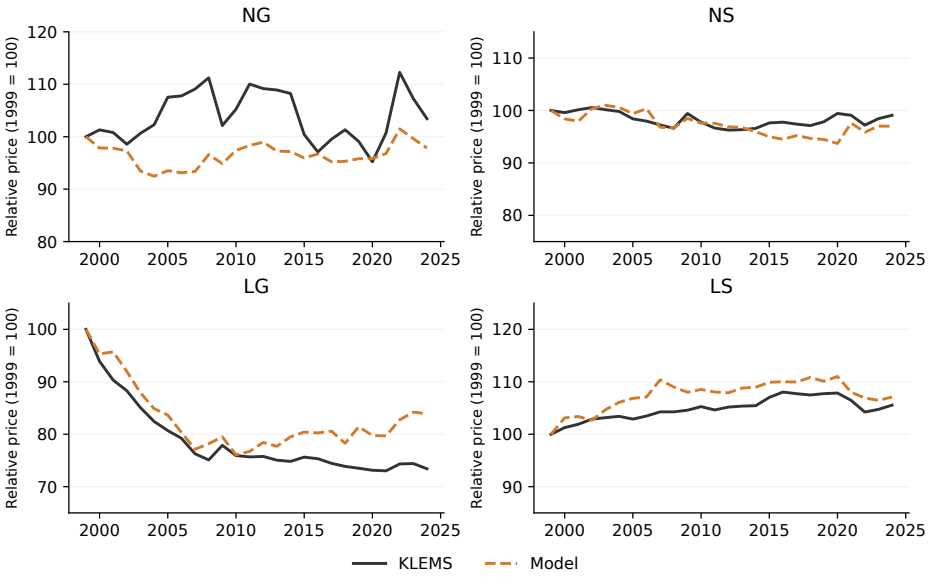


B. Hours shares

Notes: KLEMS nominal value added and CPS actual hours, 1997–2024; four classified sectors. Initial value-added shares are targets. Initial hours enter the labor-bill reconciliation; later sectoral shares are model outcomes.

Appendix E reports wages by occupation and sector over time. Worker-type wages refer to the five fixed occupation groups, while sectoral average pay also reflects changes in the mix of worker types employed.

Figure 11: Relative producer prices in the data and model



Notes: KLEMS value-added deflators and model prices, 1999–2024, relative to their respective aggregate price indexes and rebased to 100 in 1999.

5 Counterfactuals and distributional implications

Nonhomothetic demand accounts for about two-thirds of the model's increase in the luxury-service expenditure share. Together with uneven productivity growth, it raises average budget-opportunity welfare in the bottom initial expenditure decile and lowers it in the top, relative to the economy with homothetic demand and common reference-productivity growth. These differences reflect both household income and consumer prices.

5.1 Counterfactual design and sectoral reallocation

I use three counterfactuals to separate the roles of nonhomothetic demand and uneven productivity growth in sectoral reallocation. In the first counterfactual, I replace each household's expenditure function with a matched homothetic function

$$\mathcal{E}_i^H(\mathbf{p}, v) = v \, \frac{\mathcal{E}_i(\mathbf{p}, u_{i0})}{y_{i0}}. \tag{26}$$

To construct the homothetic counterfactual, I make sure that each household starts with the same consumption basket as in the nonhomothetic baseline. Here, u_{i0} denotes the household's initial utility under nonhomothetic preferences, y_{i0} its initial budget after financing, and v the homothetic utility index. At $v = y_{i0}$, the two preference systems imply the same expenditure, Hicksian quantities and compensated price responses at the initial prices. Under homothetic preferences, however, budget shares remain constant as resources increase at fixed prices, since expenditure and Hicksian quantities grow proportionally with v.

I construct this match separately for each household because households with the same taste weights can still have different consumption baskets. This is because their initial expenditure

levels put them at different points on their nonhomothetic Engel curves. I match each household separately so that both economies start from the same allocation. I can then examine how their outcomes differ when I remove the income effects of nonhomothetic demand.

The second counterfactual gives every sector the same annual rate of reference productivity growth. I construct that rate using observed midpoint nominal-value-added shares from the saved historical baseline:

$$\Delta \log T_t^{\text{equal}} = \sum_s \frac{S_{s,t-1}^{\text{data}} + S_{s,t}^{\text{data}}}{2} \Delta \log T_{st}^{\text{ref}}. \tag{27}$$

I use these annual rates to construct a common growth factor, which I apply to each sector's initial reference TFP, T_{s0}^{ref}. The weights used to calculate the common growth rate are the average observed nominal value-added shares of each sector in the current and previous year. I keep these weights fixed across all counterfactuals, instead of recalculating them from each economy's equilibrium. In contrast, the expenditure weights used for constructing the reporting price index depend on the endogenous sectoral spending shares in each economy. Each sector keeps its production-weight correction Γ_{st}^{ref}; common reference TFP growth therefore need not imply common Hicks-neutral productivity growth, realized TFP growth, or labor productivity growth (Appendix D.3). In the third counterfactual, I combine common reference-productivity growth with homothetic demand.

Across the four economies (baseline and three counterfactuals), I change preferences, reference-productivity growth, or both. Household ownership, joint-replica assignments, population weights, aggregate factor-supply paths, financing coefficients, and the remaining calibrated production paths are common across them. I keep the calibrated labor-demand shifts unchanged in all counterfactuals. In each counterfactual, I let wages, consumer prices, capital returns, household budgets, financing payments and receipts, and sectoral allocations adjust in equilibrium. As a result, household incomes and their shares of total income can also change, even under homothetic demand. Relative prices can also change under common reference-productivity growth because factor returns and sectoral input mixes continue to evolve over time.

Imposing homothetic demand reduces the 1997–2024 increase in the luxury-service expenditure share from 7.54 to 2.39 percentage points, a reduction of about two-thirds. It reduces the increase in their hours share from 8.86 to 3.90 points, or just over half.

Table 5: Counterfactual sectoral allocation, 1997–2024

Economy	LS expenditure	LS hours
Baseline	7.54	8.86
Matched homothetic	2.39	3.90
Common productivity growth	5.76	6.95
Both changes	0.57	1.86

Notes: Changes in luxury-service expenditure and hours shares are in percentage points. Common reference-productivity growth uses observed midpoint nominal-value-added weights.

Common reference-productivity growth also reduces luxury-service expansion, but less than imposing homothetic preferences. The expenditure share of luxury services still rises by 5.76 percentage points, and their hours share by 6.95 points. When both changes are imposed, these increases fall to 0.57 and 1.86 points. The combined effect is not simply the sum of the two separate effects (Table 5). This is because the effect of sectoral productivity differences depends on household preferences. Imposing homothetic demand changes how households spend additional income, which in turn affects relative prices and household incomes in equilibrium.

These counterfactuals differ from the demand accounting in Section 3. There, I separate income and relative-price effects taking the observed paths of household budgets and prices as given. Here, I change preferences or sectoral productivity growth and let both household budgets and relative prices adjust in equilibrium. That is the key difference between the historical demand accounting exercise and the model's predictions. Appendix E.2 applies the demand account to the model's own equilibrium income and price paths (Figure 45), connecting the two exercises. Hence, nonhomothetic demand can account for most of the model's luxury-service expansion even though relative-price effects contribute more than income effects in the historical account.

Appendix E.3 reports all four sectors' allocation paths. The empirical account concerns domestic consumption value added over 1999–2023, whereas the model allocates producer value added over 1997–2024, so their magnitudes also reflect different objects and periods.

5.2 Household welfare

I measure household welfare by the consumption opportunities available at each economy's prices and household budget. I value these opportunities using the same initial nonhomothetic preferences in every economy, including the homothetic counterfactual. I then express each utility level as the expenditure needed to attain it at initial prices and compare these amounts relative to the household's initial total budget.[6]

Let b denote the baseline economy and let a denote a counterfactual, and let $V_i^0(\mathbf{p}, y) = \max_{\mathbf{c}:\, \mathbf{p}'\mathbf{c} \le y} \mathcal{U}_i^0(\mathbf{c})$ be the highest utility attainable at prices $\mathbf{p}$ and budget y under initial preferences. The baseline's welfare advantage is

$$\Delta \mathrm{EV}_{it}^{b-a} = \frac{\mathcal{E}_i^0\left(\mathbf{p}_0, V_i^0(\mathbf{p}_t^b, y_{it}^b)\right) - \mathcal{E}_i^0(\mathbf{p}_0, V_i^0(\mathbf{p}_t^a, y_{it}^a))}{y_{i0}}. \tag{28}$$

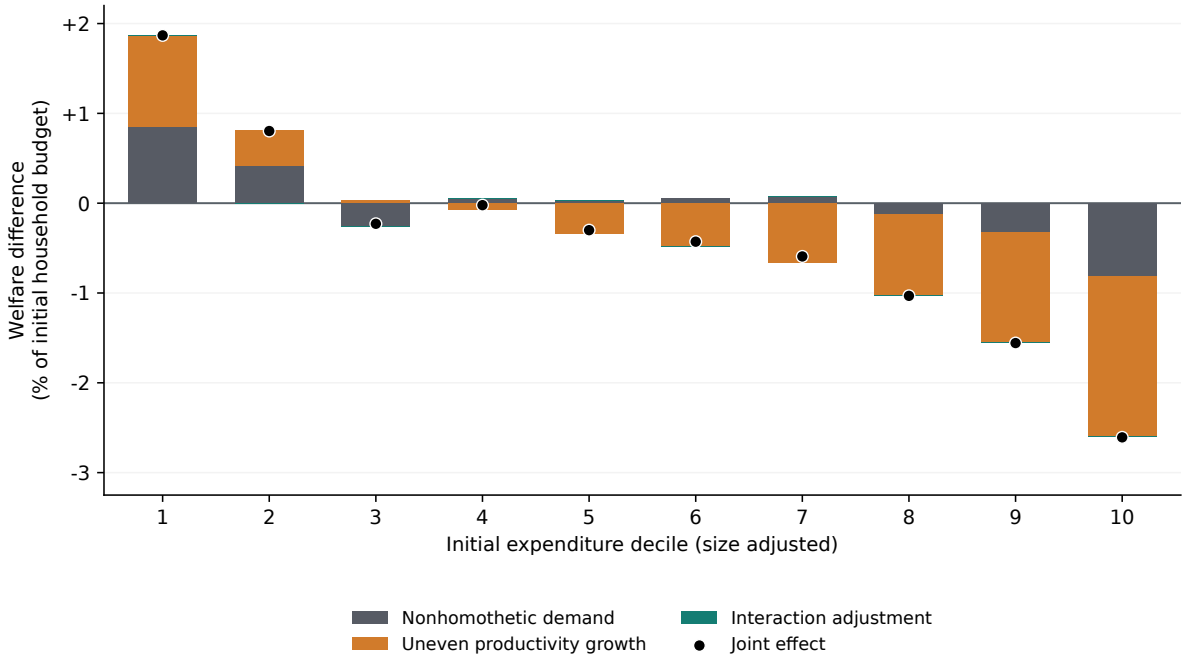
For the welfare calculation, I take each economy's equilibrium prices and household budgets, including expenditure financing, as given. I then let households reoptimize their consumption under the reference preferences. This only changes how I evaluate their welfare; I do not solve for a new equilibrium. All economies start from the same prices and household budgets. I rank all households by observed 1997 four-cell consumption-VA expenditure divided by the square root of initial household size, using fixed household-population deciles. This ranking is identical to using the proportional calibrated after-financing budget. These initial decile assignments remain fixed in every welfare figure, regardless of how household resources change across economies or

[6] This household money metric compares budget sets. It differs from the macro welfare measure in Baqaee and Burstein (2023), which compares production possibility frontiers.

over time. Positive values favor the baseline. Within each decile, I report the household-weighted mean of $100\Delta EV_{it}^{b-a}$; the denominator remains each household's initial total budget.

Nonhomothetic demand has different welfare consequences across households (gray components in Figure 12). Relative to the matched homothetic economy, the baseline yields a gain of 0.85% of initial budgets in the bottom size-adjusted expenditure decile and a loss of 0.81% in the top. The household-weighted average is close to zero (−0.004%), concealing gains and losses associated with different sources of household resources.

Figure 12: Budget-opportunity welfare by initial size-adjusted expenditure decile



Notes: Household-weighted 2024 welfare, baseline minus counterfactual, in percent of initial total budgets. Valuation uses common initial nonhomothetic preferences and prices. Fixed deciles rank observed 1997 consumption-VA expenditure divided by the square root of household size. Dots: Joint effect. Interaction: Joint effect minus the two separate comparisons.

Figure 12 reports the full general-equilibrium welfare effects of nonhomothetic demand and uneven sectoral productivity growth, with each equilibrium's prices and incomes valued under common baseline preferences. The gray and orange components correspond to the driver changed in each experiment. Each includes effects through both household incomes and consumer prices. Sectoral reallocation contributes to these welfare effects, but the comparisons do not measure its contribution separately.

Uneven productivity growth shifts gains toward households with lower initial expenditure (orange components). Relative to the common reference-growth economy, the baseline yields a gain of 1.01% of initial budgets in the bottom decile and a loss of 1.79% in the top. The black dots show the joint comparison with homothetic demand and common reference growth: The baseline yields a gain equivalent to 1.87% of initial total budgets in the bottom decile and a loss of 2.61% in the top. The joint welfare difference is therefore 4.48 percentage points larger at the bottom of the distribution.

These differences compare the baseline with the counterfactual economy in 2024; they are not cumulative gains over the period. Average budget opportunities rise from 1997 to 2024 in every initial expenditure decile under all four economies (Appendix E.4, Figure 51). Appendix Figure 52 reports the evolution of the counterfactual gaps and their equally weighted annual average over 1998–2024.

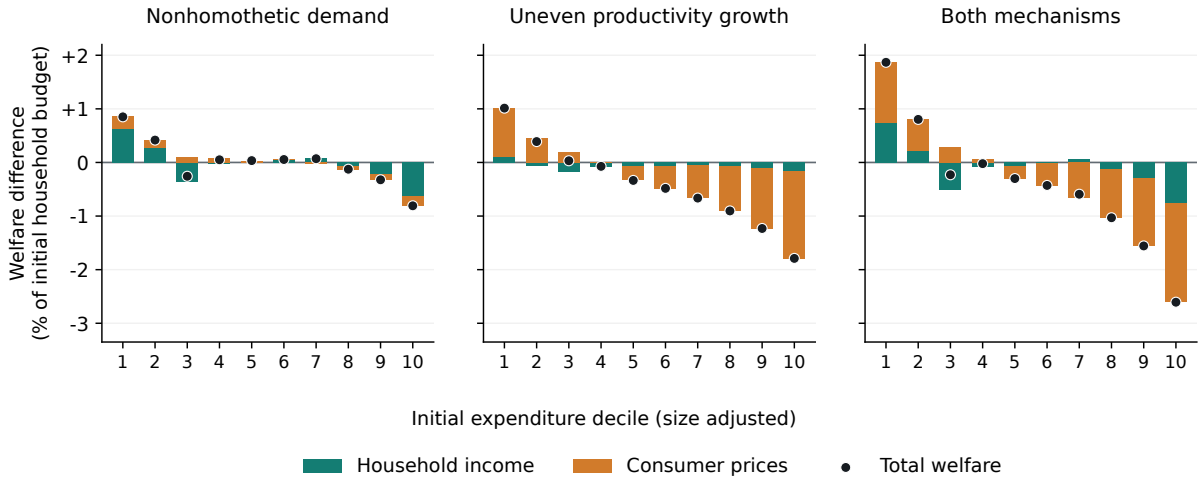
The joint welfare difference is close to the sum of the two separate comparisons. The teal component reports the remaining interaction, defined as the joint effect minus the two separate effects. It is small in every decile; at the top, it subtracts 0.01 percentage points from the sum of the separate comparisons.

5.3 Income and consumer-price channels

Both household income and consumer prices contribute to the bottom decile's gain from nonhomothetic demand, while price changes account for nearly all of its gain from uneven productivity growth. Figure 13 separates these channels for each of the three comparisons. Household income includes wages, capital returns, and expenditure financing.

For each comparison, I value baseline prices with counterfactual income and counterfactual prices with baseline income, then average the income-first and price-first changes. This splits the interaction equally between the income and price channels. For reporting, I multiply each economy's prices and incomes by the common CPI-U index divided by its own aggregate Divisia index (Appendix D.3). This normalization affects the division between income and price contributions, while leaving the total welfare difference unchanged. I evaluate the hybrid combinations of prices and incomes without solving new equilibria. Appendix D.4 gives the decomposition formulas and separates out the interaction between the two counterfactual drivers.

Figure 13: Income and consumer-price contributions by counterfactual



Notes: 2024 baseline-minus-counterfactual income and price effects, in percent of initial total budgets; valuation as in Figure 12. Teal: Income after financing. Orange: Consumer prices. Dots: Total. Contributions average income-first and price-first valuations.

For the bottom decile, the nonhomothetic-demand comparison consists of 0.63 percentage points from household income and 0.23 from consumer prices, adding to the 0.85% welfare gain before rounding. The price contribution reflects how income-driven demand changes relative production costs. As incomes rise, stronger demand for luxury services changes factor demand and equilibrium prices. Relative to the aggregate price level, luxury services are 0.93% more expensive in the baseline than in the homothetic economy, while necessity services are 1.27% cheaper. The average necessity budget share in the 2024 baseline is 53.4% in the bottom decile, compared with 38.5% in the top. This consumption mix makes the baseline price pattern more favorable to households with lower expenditure. I calculate the consumer-price contribution by changing prices while keeping household income fixed at each step.

Under uneven productivity growth, luxury goods are cheaper and luxury services more expensive than in the common-growth economy. I measure prices relative to each economy's own aggregate Divisia index (Appendix E.3, Figure 50). Richer households benefit from cheaper luxury goods, but they also lose from more expensive luxury services. The latter matters particularly for them because luxury services account for a larger share of their spending. Taking all four sectoral prices into account, the price contribution is positive in the bottom decile and negative in the top. In the bottom decile, the uneven-growth comparison contains a price contribution of 0.91 percentage points and an income contribution of 0.10. In the top decile, the corresponding contributions are -1.63 and -0.16.

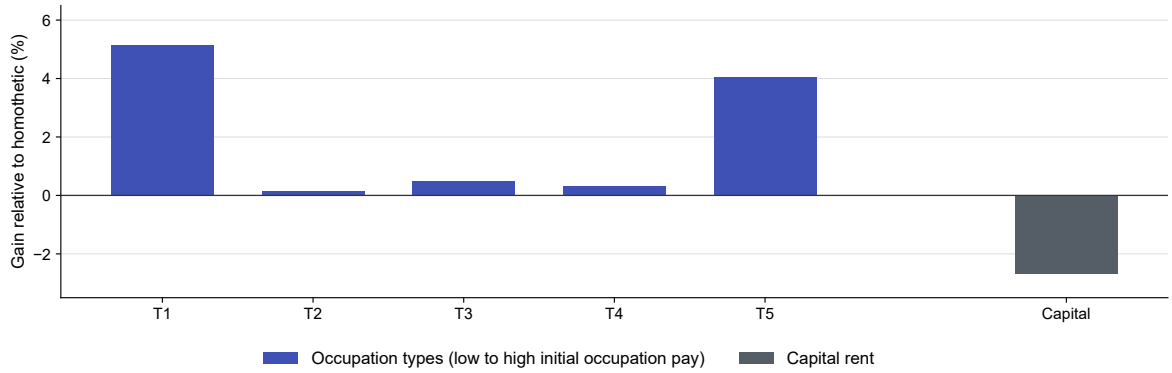
Together, the two drivers and their interaction raise welfare in the bottom decile by 0.75 percentage points through income and 1.12 percentage points through prices. In the top decile, these channels reduce welfare by 0.76 and 1.84 percentage points, respectively.

The income effects come from changes in occupation wages and capital returns. Figure 14 compares these returns in the baseline with those in the matched homothetic economy. Relative to homothetic demand, nonhomothetic demand raises real wages by 5.14% for the lowest-paid occupation group and 4.06% for the highest-paid group, but by only 0.16–0.51% for the three groups in the middle. The real rental rate of capital, by contrast, falls by 2.68%.

The wage response is U-shaped because luxury services employ both low- and high-paid occupation groups intensively. As incomes rise, households spend more on these services, which increases demand for both types of labor. This is closely related to the income-driven labor-demand channel in Comin, Danieli, et al. (2026), who also find that this mechanism leads to wage polarization. Importantly, I compare fixed occupation groups here, not households ranked by income or expenditure. The U-shaped wage response also remains when I change occupation grouping or the elasticity of substitution between labor types (Appendix E.7, Figure 63), and also when I hold the occupation-demand weights fixed at their initial values (Appendix E.8, Table 21).

Appendix Figures 48 and 49 show wage paths by occupation group and by sector. Sectoral wages also reflect the changing mix of workers in each sector. Hence, a higher average wage in luxury services does not mean that all worker types used in luxury services experience the same wage gain.

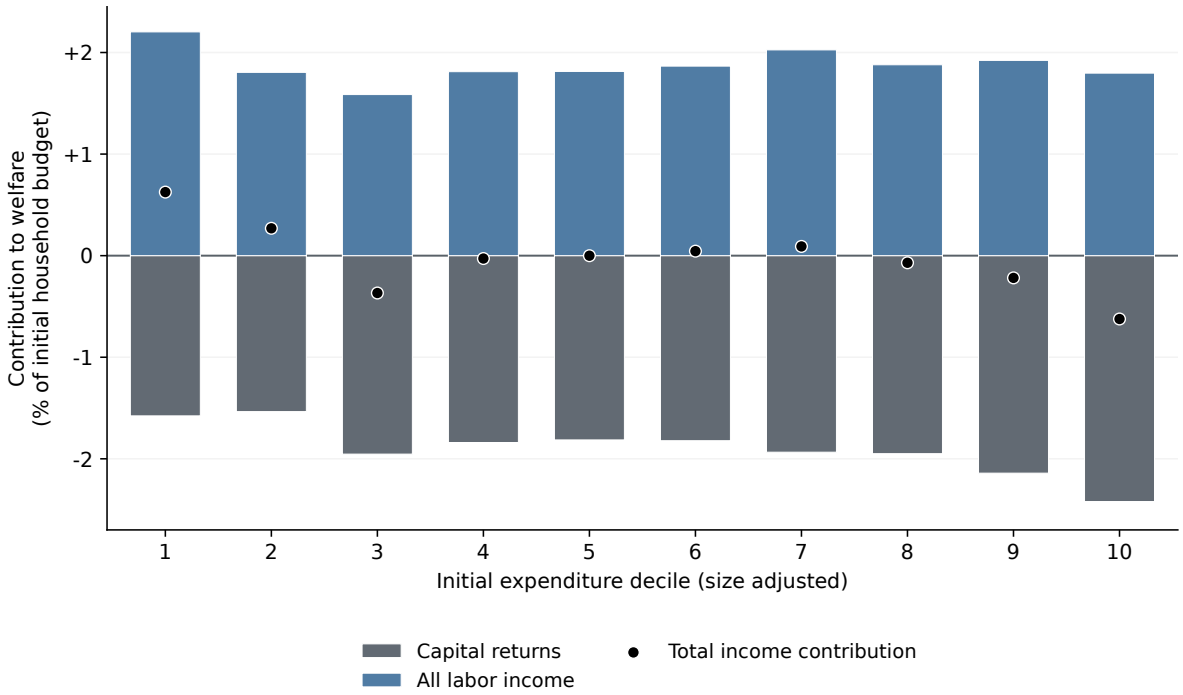
Figure 14: Nonhomothetic demand and real factor returns



Notes: 2024 real factor-price differences, 100 × (baseline/homothetic − 1). T1–T5 are fixed occupation groups, ranked by initial occupation pay across all workers. Capital is the rental rate per unit of capital services. Factor returns are deflated using each economy's own aggregate Divisia deflator.

Figure 15 links changes in factor returns to household resources by separating the income component of the nonhomothetic-demand comparison into capital and labor contributions. The labor contribution combines all five occupation types.

Figure 15: Factor-return contributions to the income channel of nonhomothetic demand



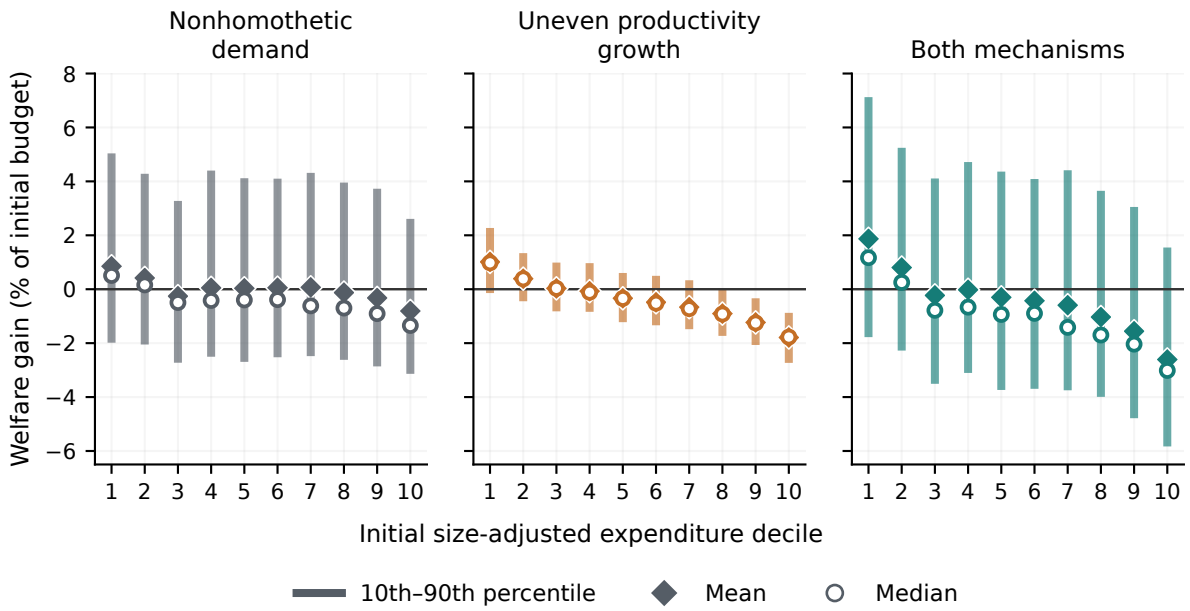
Notes: Capital and all-labor contributions to the income effect of nonhomothetic demand in Figure 13. Shapley values average the two switching orders at both economies' prices and include associated financing. Units: Percent of initial total budgets. Dots: Total income effect.

Higher labor income more than offsets the loss from lower capital returns in the bottom decile, but not in the top (Figure 15). The capital contribution is −1.58% of initial budgets in the bottom decile and −2.42% in the top. Lower capital returns thus reduce resources in both deciles, with a smaller negative contribution at the bottom. Labor contributes 2.20% in the

bottom decile and 1.80% in the top, giving net income contributions of 0.63% and −0.62%, respectively. These contributions include the financing flows, so households who hold little capital can still be affected whenever changes in capital returns (for other people) affect the payments they receive. The pattern of distributional welfare gains across households thus differs from the U-shaped wage response across occupations because households have different labor and capital holdings and financing flows.

Welfare variation within expenditure deciles. There is rich variation in welfare effects within expenditure deciles: Every expenditure decile contains households who gain and households who lose. In the joint comparison, the bottom decile gains 1.87% on average, but its 10th–90th percentile range runs from a loss of 1.78% to a gain of 7.12%. The top decile loses 2.61% on average, yet its corresponding range also includes gains, running from a loss of 5.83% to a gain of 1.55%. Figure 16 reports the mean, median, and these percentile ranges for each counterfactual. The effects of nonhomothetic demand differ widely even among households in the same decile. Initial expenditure is a better guide to the gains and losses from uneven productivity growth.

Figure 16: Welfare variation within expenditure deciles



Notes: 2024 baseline-minus-counterfactual welfare, in percent of initial total budgets. Grouping and valuation as in Figure 12. Diamonds: Means. Circles: Medians. Lines: Weighted 10th–90th percentiles across model households, not confidence intervals.

Appendix E.4 relates this variation to labor supply and other initial household characteristics (Figures 53–55). The percentile ranges here and the appendix group comparisons retain joint occupation replicas; the appendix scatterplots instead average those replicas within each original CEX household.

5.4 Which sectors drive the expansion of luxury services?

Fast productivity growth in luxury goods and slow productivity growth in luxury services both contribute to the expansion of luxury services. Starting from the historical baseline, I replace one sector's annual reference TFP growth at a time with the common growth rate in equation (27). I calculate this four-sector weighted mean using observed midpoint shares of nominal value added before changing each sector's growth path one at a time. The other three sectors meanwhile keep their historical reference TFP paths. I also keep each sector's production-weight correction Γ_{st}^{ref}, nonhomothetic preferences, and all other historical inputs the same. The resulting changes in relative prices and household purchasing power (due to equilibrium effects) then determine by how much aggregate spending shifts toward luxury services.

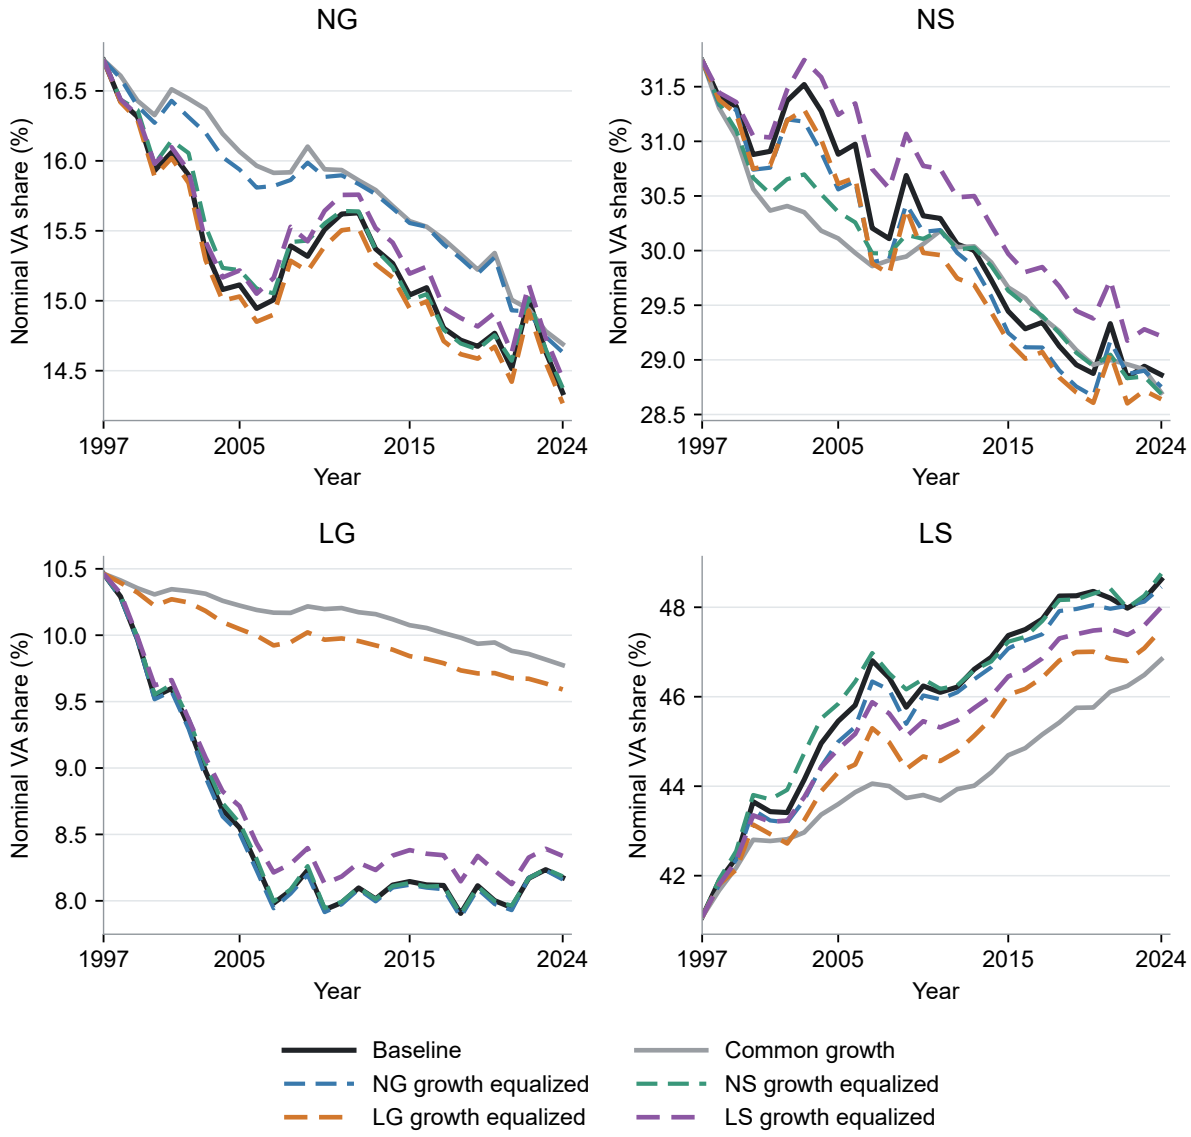
Figure 17 shows the resulting nominal value-added shares, which equal expenditure shares in the model. Equalizing luxury-goods growth reduces the luxury-service expenditure-share increase from 7.54 to 6.43 percentage points. The 1.12-point difference is 14.8% of the baseline increase. Equalizing luxury-services growth reduces the increase to 6.93 points, a difference of 0.62 points. By comparison, equalizing necessity-goods or necessity-services growth yields increases of 7.38 and 7.68 points. The fast productivity growth of luxury goods and slow productivity growth of luxury services both contribute more to this expansion than the corresponding growth patterns of the necessity sectors.

Replacing one sector's reference-productivity growth without an offset elsewhere changes both its relative growth and the observed-share-weighted mean of reference growth. Appendix E.5 also considers an experiment that adjusts growth in the other three sectors to preserve this mean in every year (Figure 57). Household purchasing power, Hicks-neutral and realized productivity, and welfare can still change.

Both luxury sectors' distinctive productivity paths support luxury-service expansion in this experiment, but their relative importance depends on the benchmark. With no offset elsewhere, equalizing luxury-goods growth reduces the increase in the luxury-service hours share by 1.23 percentage points, compared with 0.57 points when luxury-services growth is equalized. Preserving mean reference growth changes these reductions to 1.10 and 1.44 points, respectively (Appendix Table 18). Thus, luxury goods have the larger hours effect in the first comparison, while luxury services have the larger effect in the second. Both effects remain much larger than those of the necessity sectors. These mean-preserving hours comparisons also change little when occupation-demand weights are fixed: Each sector-specific reduction changes by less than 0.02 percentage points (Appendix Figure 66).

To understand what drives the increase in the luxury-service expenditure share, I use a Shapley decomposition with five components: Nonhomothetic demand and productivity growth in each of the four sectors. For demand, I compare nonhomothetic preferences with the matched homothetic benchmark. For productivity, each sector either follows its historical reference TFP growth or the same fixed common-growth path, without offsetting changes in productivity elsewhere. This yields 32 combinations, all of which I solve. I then calculate each component's contribution by averaging its marginal contribution across the 120 possible orders in which the

Figure 17: Sectoral nominal value-added shares after equalizing one sector's productivity growth



Notes: Nominal value-added shares, 1997–2024. Each colored line replaces the mentioned sector's reference TFP growth with the annual four-sector average growth rate. All other sectors meanwhile keep their historical growth.

five components can be switched. I report separately the luxury-service expansion that remains under homothetic preferences and common growth. Appendix E.5 states the formula.

Table 6: Sources of luxury-service expenditure-share expansion

Component	Contribution (pp)	Share (%)
Nonhomothetic preferences	5.17	68.53
Luxury goods productivity	1.00	13.20
Luxury services productivity	0.77	10.25
Necessity goods productivity	0.14	1.86
Necessity services productivity	-0.11	-1.46
Homothetic demand and common growth	0.57	7.62
Total	7.54	100.00

Notes: Shapley contributions to the 1997–2024 increase in the luxury-service expenditure share, in percentage points. I calculate each contribution by averaging the component's marginal effect across all 120 possible orders in which I change preferences and productivity growth in the four sectors. For each sector, I compare historical reference TFP growth with the same four-sector weighted average. I do not adjust growth in the other sectors to offset this change (Appendix E.5).

Table 6 shows that nonhomothetic preferences account for 5.17 percentage points, or 68.5%, of the total increase in the luxury-service expenditure share. Relatively strong productivity growth in luxury goods contributes another 1.00 percentage point (13.2%), while relatively weak productivity growth in luxury services contributes 0.77 points (10.3%). The contributions of necessity goods and services are small and partly offset each other. Together, differences in productivity growth across the four sectors account for 1.80 percentage points of the increase, of which 55.3% comes from luxury goods and 43.0% from luxury services. The Shapley contribution of luxury goods is slightly smaller than the 1.12-point effect of equalizing its productivity growth alone. This is because the Shapley decomposition averages this effect across different combinations of preferences and productivity growth in the other sectors.

Nonhomothetic demand is the largest contributor in every switching order. Its contribution ranges from 4.85 to 5.49 percentage points, compared with at most 1.13 points for any individual productivity component (Appendix Table 19).

The remaining 0.57 percentage points (7.6%) arise under matched homothetic preferences and common reference-productivity growth. Household responses to changing relative prices account for 0.45 points, and changes in households' shares of aggregate spending account for 0.12 points (Appendix E.5, equation (68)). Relative prices still change as factor supplies, production weights, and factor returns change over time. Households in the matched homothetic economy also still have different consumption baskets from each other. Hence aggregate expenditure shares can change as some households account for more of total spending, even though each household's budget shares do not change with its own income at fixed prices.

Growth accounting shows that luxury services account for most of the productivity drag from sectoral reallocation. The model helps explain why these services expand in the first place: The rapid productivity growth in luxury goods expands luxury services, even more so than the weak productivity growth in luxury services itself. In the Shapley decomposition, these two luxury sectors make up almost the entire contribution of uneven productivity growth. Most of the remaining increase in the luxury-service expenditure share is due to nonhomothetic demand.

These results also show why it is important to distinguish necessities from luxuries within both goods and services: A simple luxury/necessity split would group together luxury goods and luxury services, even though their productivity growth differs substantially. It would therefore hide how productivity growth in both sectors contributes to the expansion of luxury services. A goods/services split instead would group luxuries with necessities, even though changes in the luxury sectors contribute much more to this expansion than changes in their necessity counterparts. By distinguishing all four sectors, I can capture both of these differences.

5.5 Robustness

The distributional pattern holds with five or ten occupation groups and labor substitution elasticities of one or 1.5. Across the three main welfare comparisons, mean welfare effects by expenditure decile change by at most 0.20 percentage points, and the occupational U-shape remains (Appendix E.7, Figures 62–63). For each grouping, I reassign workers to occupation types based on their initial pay, while keeping the adjustment for changes in CPS occupation codes over time, the aggregate labor-services target from KLEMS and household capital holdings unchanged.

The results are also robust to holding occupation-demand weights fixed at their initial values, rather than adjusting them to match historical relative wage growth. For this exercise, I recompute the reference-productivity correction using these fixed weights. The model then fits historical relative wages much less well, but the distributional results change little: Mean welfare effects by expenditure decile differ by at most 0.32 percentage points, and the U-shaped wage response remains (Appendix E.8, Figures 64–65 and Table 21).

Finally, I consider an alternative rule for financing household expenditure. The main distributional pattern remains: Poorer households still gain on average, while richer households lose (Appendix E.6, Figure 60). In this exercise, I keep the same households and their initial budgets as in the baseline. For households with positive initial factor income, I let expenditure grow at the same rate as their own factor income, with a common adjustment to ensure that aggregate household budgets add up to aggregate income. For households without initial factor income, I instead keep their initial share of aggregate income fixed. I then solve all four economies again under this alternative rule. The results are similar overall, although some small gains from non-homothetic demand in the middle deciles become losses. The largest change in average welfare across deciles is 0.62 percentage points.[7]

Across these robustness checks, poorer households always gain on average, while richer households always lose. The results for the middle deciles depend a bit more on how expenditure is financed (through factor incomes only or also through transfers). In all counterfactuals, I keep aggregate factor-supply paths, household ownership, and occupation assignments the same. This means that saving, labor supply and occupation assignments do not respond. The welfare comparisons therefore capture how the demand and productivity forces behind sectoral reallocation

[7] Appendix Figure 61 removes financing as a broader stress test. It compares the same households with positive factor income, covering 98.04% of the original population weight, but also changes the initial equilibrium and spending fit. I rematch the homothetic benchmark to that allocation.

affect households through changes in their incomes and consumer prices, rather than the effects of reallocation alone.

6 Conclusion

Economic growth changes not only how much households consume, but also which sectors expand and who benefits from their expansion. This paper shows that these changes can pull in different directions: The expansion of luxury services accounts for most of the measured drag on U.S. labor productivity growth from sectoral reallocation, while the demand and productivity forces behind that expansion disproportionately benefit poorer households.

Over 1980–2024, observed labor productivity growth was 0.20 percentage points per year below the fixed-composition accounting benchmark, with luxury services accounting for about two-thirds of this gap. While both luxury goods and luxury services attract a larger share of household spending at higher expenditure levels, they differ substantially in how they are produced. Luxury goods have experienced rapid productivity growth and declining domestic employment, while luxury services have weak productivity growth and a growing share of employment. Distinguishing luxuries from necessities within both goods and services therefore reveals a dimension of structural change that neither split alone captures.

Why have luxury services expanded? I find that both rising household expenditure and relative-price changes contribute substantially. As households become richer, they devote larger shares of their budgets to luxury services. At the same time, these services have become more expensive relative to other sectors. Because households substitute away from them only to a limited extent, their expenditure share rises as their relative price increases. Surprisingly, part of this expansion originates in strong productivity growth in luxury goods, which raises household purchasing power and makes luxury services relatively more expensive. In my model, nonhomothetic demand accounts for about 69% of the increase in the luxury-service expenditure share, while relatively strong productivity growth in luxury goods accounts for another 13%. Productivity growth in one luxury sector thus helps expand another luxury sector in which productivity growth is much weaker.

The same forces that expand luxury services also have very different consequences for poorer and richer households. Luxury services are labor-intensive and use both low- and high-paid occupations intensively. In my model, relative to homothetic demand, nonhomothetic demand raises real wages most strongly for these occupation groups while lowering the real rental rate of capital. Poorer households benefit from stronger labor demand and are less exposed to lower capital returns. They also spend smaller budget shares on luxury services and are therefore less affected by the increase in their relative price. Richer households rely more on capital income and devote larger budget shares to these services. Relative to an initially matched economy with homothetic demand and common sectoral productivity growth, the baseline yields average welfare gains in 2024 equivalent to 1.9% of initial budgets for households in the bottom expenditure decile, compared with losses of 2.6% for those in the top decile. These differences reflect the combined effects of demand and sectoral productivity growth, rather than the welfare

consequences of reallocation alone. They also conceal substantial heterogeneity within every expenditure decile.

An interesting application of my framework is to study how artificial intelligence (AI) might affect the luxury shift. Luxury services involve many non-routine and interpersonal tasks (Appendix Figure 25), so their productivity response to AI may differ from that of other sectors. If AI raises productivity mainly outside of luxury services, the resulting income growth could increase spending on luxury services (as they are the most expenditure-elastic sector), while their relative price would increase even more. With limited substitution away from them, this could further increase their expenditure share and draw more workers into the sector. If AI instead raises productivity faster in luxury services than elsewhere, these services could become relatively cheaper, and the shift in employment toward them could slow or stop. My model could be used to quantify how these different productivity paths affect aggregate productivity growth, occupational wages and household welfare.

Existing work on AI and inequality emphasizes its effects on labor earnings and capital income (Acemoglu 2025; Rockall, Mendes Tavares, and Pizzinelli 2025). My framework could examine a counteracting force: Higher incomes from AI could lead households to spend a larger share of their budgets on labor-intensive luxury services. This shift might put downward pressure on capital returns and partly offset the increase predicted by Rockall et al. (2025). The distributional consequences would also depend on changes in relative consumer prices, since poorer and richer households consume different baskets. Arvai and Mann (2026) show that the price effects of digitalization disproportionately benefit richer households. My framework could be used to examine whether AI-driven reallocation toward luxury services supports poorer households' labor earnings while raising the relative prices of services on which richer households spend larger budget shares.

More generally, rising incomes can continue to shift spending toward luxury services. If productivity grows more slowly in these services than elsewhere and substitution to other sectors remains limited, this shift could remain a drag on aggregate productivity growth. Yet this drag alone does not reveal who gains from the forces behind the shift. This paper shows that those forces can disproportionately benefit poorer households.

References

Acemoglu, Daron. 2025. "The Simple Macroeconomics of AI." *Economic Policy* 40 (121): 13–58.

Aguiar, Mark, and Mark Bils. 2015. "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review* 105 (9): 2725–2756.

Andreolli, Michele, Natalie Rickard, and Paolo Surico. 2025. "Non-Essential Business Cycles." Discussion Paper 19773. CEPR.

Arvai, Kai, and Katja Mann. 2026. "Consumption Inequality in the Digital Age." *American Economic Journal: Macroeconomics*. Forthcoming.

Baqaee, David R., and Ariel Burstein. 2023. "Welfare and Output with Income Effects and Taste Shocks." *The Quarterly Journal of Economics* 138 (2): 769–834.

Baumol, William J. 1967. "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis." *American Economic Review* 57 (3): 415–426.

Bergman, Nittai K., Nir Jaimovich, and Itay Saporta-Eksten. 2026. "The Impact of Economic Shifts: From Micro-Elasticities to Distributional Welfare Effects." Working Paper.

Boppart, Timo. 2014. "Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences." *Econometrica* 82 (6): 2167–2196.

Comin, Diego, Ana Danieli, and Martí Mestieri. 2026. "The Income Channel to Labor-Market Polarization." Working Paper 35315. National Bureau of Economic Research.

Comin, Diego, Danial Lashkari, and Martí Mestieri. 2021. "Structural Change with Long-Run Income and Price Effects." *Econometrica* 89 (1): 311–374.

Comin, Diego, Danial Lashkari, and Martí Mestieri. 2025. "The Structural Transformation of Innovation." Working Paper 33855. National Bureau of Economic Research.

Cravino, Javier, and Sebastian Sotelo. 2019. "Trade-Induced Structural Change and the Skill Premium." *American Economic Journal: Macroeconomics* 11 (3): 289–326.

Deaton, Angus, and John Muellbauer. 1980. "An Almost Ideal Demand System." *American Economic Review* 70 (3): 312–326.

Duernecker, Georg, Berthold Herrendorf, and Ákos Valentinyi. 2024. "Structural Change within the Services Sector and the Future of Cost Disease." *Journal of the European Economic Association* 22 (1): 428–473.

Fan, Tianyu, Michael Peters, and Fabrizio Zilibotti. 2023. "Growing Like India—The Unequal Effects of Service-Led Growth." *Econometrica* 91 (4): 1457–1494.

Fleck, Johannes, Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante. 2025. "Fiscal Progressivity of the U.S. Federal and State Governments." Working Paper 33385. National Bureau of Economic Research.

Foellmi, Reto, and Josef Zweimüller. 2006. "Income Distribution and Demand-Induced Innovations." *The Review of Economic Studies* 73 (4): 941–960.

Foellmi, Reto, and Josef Zweimüller. 2008. "Structural Change, Engel's Consumption Cycles and Kaldor's Facts of Economic Growth." *Journal of Monetary Economics* 55 (7): 1317–1328.

Hall, Robert E., and Charles I. Jones. 2007. "The Value of Life and the Rise in Health Spending." *The Quarterly Journal of Economics* 122 (1): 39–72.

Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi. 2013. "Two Perspectives on Preferences and Structural Transformation." *American Economic Review* 103 (7): 2752–2789.

Ho, Paul, Danial Lashkari, and Pierre-Daniel G. Sarte. 2026. "How Do Sectoral Shocks Shape Future GDP?" Working Paper 26-11. Federal Reserve Bank of Richmond.

Hubmer, Joachim. 2023. "The Race Between Preferences and Technology." *Econometrica* 91 (1): 227–261.

Jaravel, Xavier. 2026. "Distributional Consumer Price Indices." Working Paper.

Jaravel, Xavier, and Danial Lashkari. 2024. "Measuring Growth in Consumer Welfare with Income-Dependent Preferences: Nonparametric Methods and Estimates for the United States." *The Quarterly Journal of Economics* 139 (1): 477–532.

Mazzolari, Francesca, and Giuseppe Ragusa. 2013. "Spillovers from High-Skill Consumption to Low-Skill Labor Markets." *The Review of Economics and Statistics* 95 (1): 74–86.

Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo. 2022. "Uneven Growth: Automation's Impact on Income and Wealth Inequality." *Econometrica* 90 (6): 2645–2683.

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

Ngai, L. Rachel, and Orhun Sevinc. 2025. "A Multisector Perspective on Wage Stagnation." *Review of Economic Dynamics* 56: 101269.

Olivi, Alan, Vincent Sterk, and Dajana Xhani. 2026. "Optimal Monetary Policy during a Cost-of-Living Crisis." *American Economic Review*. Forthcoming.

Rockall, Emma J., Marina Mendes Tavares, and Carlo Pizzinelli. 2025. "AI Adoption and Inequality." Working Paper 2025/068. International Monetary Fund.

Wiedemann, Lukas. 2026. "The Dual Role of Inequality for Structural Change: Cause and Consequence." Working Paper.

A Data and sector classification

A.1 Category-level Engel estimates

The classification uses fixed expenditure elasticities for 35 consumption categories, estimated from the 1980–2020 CEX Interview Survey following Aguiar and Bils (2015). The category estimates require complete panels of four expenditure interviews, standardized to interviews 2–5 across survey designs. The sample consists of urban consumer units whose reference person is aged 25–64. I retain the original income and expenditure screens: Nonnegative reported after-tax income, positive food expenditure below total expenditure, and before-tax income within the year's 5th–95th percentile range. Expenditure in any of the 16 original categories other than food at home, food away from home, housing, and vehicles cannot exceed half of after-tax income. Final survey weights are normalized within year.

For each consumption category c, I estimate the following Engel curve:

$$\widetilde{x}_{hct} \equiv \frac{x_{hct} - \bar{x}_{ct}}{\bar{x}_{ct}} = \alpha_{ct} + \eta_c \ln X_{ht} + \boldsymbol{\gamma}_c' \mathbf{Z}_{ht} + u_{hct}. \tag{29}$$

Here h indexes households and t years. The variable x_{hct} is household h's annual real expenditure on category c, and $\bar{x}_{ct}$ is its survey-weighted mean across households in year t. Total real household expenditure, X_{ht}, retains the original 20-category Aguiar–Bils aggregate. The category-specific year effects α_{ct} absorb changes common to households within each category and year. The vector $\mathbf{Z}_{ht}$ contains indicators for the reference person's age (38–50 and 51–64, relative to 25–37), household size (3–4 and five or more, relative to one or two), and two or more earners. The coefficients $\boldsymbol{\gamma}_c$ vary across categories, and u_{hct} contains unobserved demand shifters and reporting error.

I express category expenditure as a proportional deviation from its annual mean so that households reporting zero spending remain in the sample. For positive expenditure, this outcome is a first-order approximation to $\ln x_{hct} - \ln \bar{x}_{ct}$ around $x_{hct} = \bar{x}_{ct}$. Hence, η_c has the interpretation of a local expenditure elasticity at the category-year mean. The dependent variable differs from the household's budget share, x_{hct}/X_{ht}.

Errors in reporting individual purchases enter both category and total expenditure, which can correlate $\ln X_{ht}$ with u_{hct}. I therefore estimate equation (29) by survey-weighted two-stage least squares. Specification I, the baseline, instruments log annual total expenditure with before-tax-income-group indicators and log real after-tax income. Specification II constructs category and total expenditure from the final two expenditure interviews and instruments log total expenditure with its counterpart from the first two interviews. Under the standardized interview numbering, these are interviews 4–5 and 2–3, respectively; the category-year mean in Specification II is also calculated from interviews 4–5. Both specifications include the demographic controls and year indicators in the first stage. The IV interpretation requires the excluded instruments to be uncorrelated with u_{hct} conditional on these controls. These category-level estimates are distinct from the demand-system estimates in Section 3.2.

Table 7 reports pooled 1980–2020 Engel estimates and their standard errors for both specifications.[8] Specification I supplies the baseline classification, and Specification II the alternative estimates. An elasticity above one implies luxury status and an elasticity below one necessity status; unit elasticities and unsupported assignments remain unclassified. The baseline assigns primary-residence housing to necessities and healthcare to luxuries, as discussed in Section 2.1. These category assignments, marked in the table, leave the estimated numerical coefficients unchanged. The comparison excluding housing and healthcare retains the same Engel specification and industry-mapping method.

Primary-residence housing has expenditure elasticities of 0.95 under Specification I and 0.94 under Specification II, placing the combined category among necessities in both specifications. When I estimate separate Engel curves, rented housing is a necessity and owner-occupied housing a luxury. However, I cannot map these categories to distinct producing industries, because both forms of housing are supplied by the same construction and real-estate industries at the four-digit NAICS level. I therefore classify housing as a necessity, consistent with the combined-housing estimates. This is also in line with Fleck, Heathcote, Storesletten, and Violante (2025), who show using American Community Survey (ACS) data that both rents and home values rise less than proportionally with household income.

The comparison excluding housing and healthcare changes the coverage of the baseline classification rather than excluding whole industries from the analysis sample.

The early and late point estimates place 28 of the 35 categories on the same side of one (Table 8). Two move from below to above one and five move in the opposite direction. These comparisons describe the stability of estimated labels. A crossing of one is not a formal rejection of a stable elasticity, and unchanged category labels do not by themselves validate every historical industry mapping.

A.2 From consumption categories to producing industries

The reviewed consumption-category-to-industry mapping supplies direct assignments to four-digit producing industries. For industries without a direct assignment, I trace exposure to household consumption through a fixed input-output network, preserving supported direct labels. The construction uses the detailed 2007 BEA Use and Make tables before redefinitions and the corresponding Import Matrix. Make-table market shares allocate commodity demand to producing industries. Imported final purchases and imported intermediate inputs are removed separately. The resulting domestic input requirements trace both direct production and upstream suppliers.

The final-demand seed is 2007 commodity personal consumption expenditure. I split each commodity across consumption categories using CEX expenditure shares averaged equally over 2005–2007 and the dated consumption-item correspondences. Unsupported category assignments remain unresolved consumption. Let r_{jc} denote domestic production in industry j required to

[8] The table uses the later saved pooled output, whose estimates differ by less than 0.005 from the original pooled coefficients retained in the industry classification.

Table 7: Consumption categories, Engel estimates, and baseline assignments

Consumption category	I (SE)	II (SE)	Baseline
Ship fares, school bus, and other transportation (residual)	2.38 (0.12)	2.22 (0.15)	L
Hotels and other lodging away from home	2.05 (0.03)	1.86 (0.04)	L
Airline fares	1.96 (0.04)	1.82 (0.05)	L
Education	1.85 (0.05)	2.07 (0.06)	L
Entertainment fees, admissions, reading	1.78 (0.02)	1.74 (0.02)	L
Domestic services and childcare	1.73 (0.04)	1.82 (0.04)	L
Cash contributions (not for alimony/support)	1.71 (0.05)	1.49 (0.04)	L
Intercity bus and train fares	1.57 (0.06)	1.45 (0.08)	L
Recreation equipment, vehicles, hobbies, fees and rentals	1.51 (0.04)	1.43 (0.05)	L
Furniture and fixtures	1.50 (0.03)	1.49 (0.04)	L
Pet and veterinary services	1.47 (0.04)	1.58 (0.05)	L
Alcoholic beverages away from home	1.47 (0.03)	1.45 (0.03)	L
Other household goods, repairs, rentals, and equipment (residual)	1.46 (0.03)	1.58 (0.04)	L
Computers, communications equipment, and software	1.39 (0.04)	1.34 (0.05)	L
Men's and women's clothing	1.29 (0.02)	1.32 (0.02)	L
Food away from home	1.21 (0.01)	1.19 (0.02)	L
Tools, garden equipment and other durable household equipment	1.21 (0.05)	1.19 (0.06)	L
Personal care products and services	1.11 (0.01)	1.12 (0.01)	L
Audio/video electronics hardware	1.10 (0.04)	1.06 (0.05)	L
Shoes and other apparel	1.08 (0.02)	1.12 (0.03)	L
Household appliances	1.08 (0.03)	1.12 (0.04)	L
Healthcare incl. insurance	0.90 (0.01)	1.03 (0.01)	L†
Vehicles incl. purchasing, leasing and insurance	0.99 (0.02)	0.79 (0.03)	N
Alcoholic beverages for home use	0.96 (0.02)	0.99 (0.02)	N
Primary-residence housing, including owner-equivalent rent	0.95 (0.01)	0.94 (0.01)	N†
Pet purchases, supplies and medicine	0.90 (0.04)	1.13 (0.04)	N
Recorded/digital media and A/V rentals and repairs	0.83 (0.02)	0.95 (0.03)	N
Children's clothing (up to age 15)	0.66 (0.02)	0.77 (0.02)	N
Vehicle fuel, parts, maintenance, and repairs	0.64 (0.01)	0.70 (0.01)	N
Local transit and taxis	0.56 (0.05)	0.44 (0.05)	N
Cable and satellite television services	0.52 (0.01)	0.53 (0.01)	N
Utilities, fuels, and public services	0.47 (0.01)	0.53 (0.01)	N
Telephone, cellular, and internet services	0.47 (0.01)	0.58 (0.01)	N
Food at home	0.39 (0.00)	0.47 (0.01)	N
Tobacco products and smoking supplies	-0.49 (0.02)	-0.12 (0.02)	N

Notes: CEX, 1980–2020; pooled IV specifications following Aguiar and Bils (2015). Baseline: Specification I. Standard errors in parentheses. N and L denote necessities and luxuries; † marks the housing and healthcare overrides.

Table 8: Classification stability across estimation periods

Early classification	Late necessity	Late luxury
Necessity in 1984–1986	11	2
Luxury in 1984–1986	5	17

Category	1984–1986	2017–2019	1980–2020
Shoes and other apparel	1.11	0.81	1.08
Personal care products and services	0.93	1.10	1.11
Vehicles incl. purchasing, leasing and insurance	1.23	0.90	0.99
Alcoholic beverages for home use	0.94	1.01	0.96
Audio/video electronics hardware	1.18	0.83	1.10
Recorded/digital media and A/V rentals and repairs	1.05	0.76	0.83
Pet purchases, supplies and medicine	1.34	0.97	0.90

Notes: CEX; Specification I for 35 categories. Point estimates above (below) one define luxuries (necessities), without housing or healthcare overrides. Lower panel: Categories crossing one between 1984–1986 and 2017–2019, with pooled 1980–2020 estimates.

supply consumption category c, including indirect requirements, and let $\mathcal{C}^*$ contain categories with usable Engel coefficients and an eligible classification. The upstream score and its support share are

$$\eta_j^U = \frac{\sum_{c\in\mathcal{C}^*} r_{jc}\eta_c}{\sum_{c\in\mathcal{C}^*} r_{jc}}, \qquad\qquad q_j = \frac{\sum_{c\in\mathcal{C}^*} r_{jc}}{\sum_c r_{jc}}.$$

Here η_c is the numerical coefficient from the selected Engel specification, including for housing and healthcare. The denominator of q_j includes all consumption-supported production, including unresolved consumption. Investment, government purchases, and exports are outside this denominator.

An upstream assignment requires $q_j \geq 0.25$. Among eligible industries, $\eta_j^U > 1$ implies luxury status and $\eta_j^U < 1$ necessity status. Industries with insufficient support, no usable coefficient, or a unit score remain unclassified. Thus, the threshold concerns the supported share of consumption-related production, not 25% of total industry sales. The score is a production-weighted mean of consumption elasticities, not an estimated elasticity of total industry revenue. I hold the network and classification fixed over time; outcome weights can vary by year.

Industry codes are harmonized to each outcome source's native vintage. BDS and the current KLEMS mapping use NAICS2017; consumer-item applications use dated correspondences. Within broader industries, I retain necessity, luxury, and unclassified child mass and construct four-sector exposures from the joint assignments. Multiplying separate luxury and goods shares would incorrectly impose independence between the classifications.

Positive activity weights allocate KLEMS and BEA parent-industry outcomes to their children. The baseline uses the most detailed supported nominal-value-added information, with BDS employment allocating BEA categories that are broader than four-digit industries. These allocations measure exposure rather than independently observed child-industry value added.

Classification weights determine sector assignments; outcome weights then follow the quantities being aggregated.

For production-side outcomes, I split additive nominal accounts, factor payments, and hours by the joint sector exposures and sum them before taking ratios. Quantity indexes instead chain industry log growth using adjacent-year midpoint shares of the corresponding nominal output, compensation, or input-cost account. Sectoral labor shares divide allocated labor compensation by allocated nominal value added. Real value added per hour divides the sectoral value-added quantity index by summed allocated hours; it is not an average of separately normalized industry productivity levels. I do not sum nonadditive chained-dollar quantities. The separate growth account in Appendix B.7 specifies its treatment of changing classified exposure.

The headline consumer-price indexes combine D-CPI price changes with CEX expenditure shares using monthly Törnqvist weights. Direct consumption-item classifications include owner-equivalent rent. Arithmetic annual averages of the monthly indexes are normalized to 100 in 1999. These final-consumption baskets differ from the producing-industry coverage of BDS and KLEMS, and their consumer prices are distinct from KLEMS value-added deflators.

Domestic value added embodied in household consumption is a third object, used in Section 3.2. Its fixed 2012 input-output bridge allocates final purchases to domestic producing origins after removing imported final purchases and imported intermediate inputs. This bridge is separate from the fixed 2007 network used for upstream classification. The resulting consumption value added differs from final household expenditure, which includes import content, and from total industry value added, which also serves other final uses. It is a fixed-technology accounting exposure, not an official GDP or welfare measure.

A.3 Classification coverage

The baseline classifies 297 of the 303 nongovernment four-digit NAICS2017 industries (Table 9). The direct mapping assigns 107 industries, and upstream completion adds 190. In 2023, the classification covers 99.07% of BDS employment and 98.74% of private KLEMS nominal value added. The different industry counts reflect each source's coverage and the positive activity weights used to allocate KLEMS industries.

Table 9: Coverage of the baseline industry classification

Universe or activity	Total	Direct	Upstream	Classified	Unclassified
A. Four-digit industry counts					
NAICS2017 universe	303	107	190	297	6
BDS observed industries	288	106	177	283	5
Positive-weight KLEMS children	301	106	189	295	6
B. Shares of activity (percent)					
BDS employment	100.00	46.40	52.67	99.07	0.93
Private KLEMS nominal VA	100.00	44.36	54.39	98.74	1.26

Notes: BDS and KLEMS, 2023; non-government activity. Panel A: Four-digit NAICS2017 industry counts; KLEMS counts only positive-weight mapped industries. Panel B: Employment and private nominal-value-added shares. KLEMS accounts use nominal-value-added classification weights. Denominators include unclassified activity.

I retain unclassified exposure when allocating broader industries, rather than distributing it among the classified cells. The four-sector share figures condition on classified activity, while the coverage table includes unclassified activity in its denominators. The main growth account excludes unclassified activity.

A.4 Luxury shares across broad industries

Table 10 shows substantial variation in luxury shares across broad industries. Accommodation and food services, education, and arts and recreation are fully classified as luxury. The luxury share is 97.9% in healthcare and social assistance and 83.9% in professional, scientific, and technical services, compared with 16.4% in real estate and rental and leasing. Unclassified employment remains in each denominator.

Table 10: Luxury employment shares within broad industries

NAICS	Broad industry	Luxury share (%)
11	Agriculture, Forestry, Fishing and Hunting	30.1
21	Mining, Quarrying, and Oil and Gas Extraction	7.1
22	Utilities	0.0
23	Construction	0.0
31	Manufacturing	11.9
32	Manufacturing	71.0
33	Manufacturing	55.3
42	Wholesale Trade	50.6
44	Retail Trade	42.3
45	Retail Trade	46.1
48	Transportation and Warehousing	92.7
49	Transportation and Warehousing	43.0
51	Information	72.4
52	Finance and Insurance	48.5
53	Real Estate and Rental and Leasing	16.4
54	Professional, Scientific, and Technical Services	83.9
55	Management of Companies and Enterprises	100.0
56	Administrative and Support and Waste Management and Remediation Services	77.8
61	Educational Services	100.0
62	Health Care and Social Assistance	97.9
71	Arts, Entertainment, and Recreation	100.0
72	Accommodation and Food Services	100.0
81	Other Services (except Public Administration)	80.0

Notes: Share of 2023 BDS employment classified as luxury, in percent, using baseline four-digit NAICS2017 assignments. Unclassified employment remains in the denominator.

A.5 Worker samples and wage measurement

Workers in broad CPS industries contribute fractionally to sectors according to the joint luxury and goods/services assignments of their constituent four-digit industries. Unclassified exposure remains in the account. These are industry exposures rather than observed four-digit assignments for individual workers. The industry-skill comparison fixes the allocation using 2007 BDS employment weights.

Table 11: Largest industries under the baseline classification

Group	NAICS	Industry	Employment
Luxury	7225	Restaurants and Other Eating Places	11.41
	5613	Employment Services	6.23
	6221	General Medical and Surgical Hospitals	5.71
	5511	Management of Companies and Enterprises	3.78
	6211	Offices of Physicians	2.84
Necessity	4451	Grocery Stores	2.96
	4523	General Merchandise Stores, including Warehouse Clubs and Supercenters	2.48
	2382	Building Equipment Contractors	2.37
	5617	Services to Buildings and Dwellings	2.16
	5413	Architectural, Engineering, and Related Services	1.67
Unclassified	5239	Other Financial Investment Activities	0.65
	6242	Community Food and Housing, and Emergency and Other Relief Services	0.25
	6243	Vocational Rehabilitation Services	0.23
	8122	Death Care Services	0.14
	5259	Other Investment Pools and Funds	0.01

Notes: BDS, 2023; the five largest four-digit NAICS2017 industries in each group by private employment, in millions.

The industry-skill split uses pooled Basic Monthly private wage-and-salary employment in 2003–2005. Industries with a college share above 50% form the higher-college-share group; the remaining industries form the lower group. I fix these assignments for 2003–2024. Employment is measured with Basic Monthly person weights and annualized by averaging the twelve monthly totals. ORG hourly wages divide weekly earnings by valid usual main-job hours and use earnings weights. Real wages use CPI-U. The samples include wage-and-salary workers aged 16–64.

The wage adjustment in Table 12 holds each sector–industry group's initial demographic composition fixed. I estimate annual arithmetic-wage regressions on a flexible age profile, sex, and three non-college education categories: Less than high school, high school, and some college or an associate degree. Predictions are averaged over the corresponding group's 2003 composition. The common weekly-earnings ceiling is held constant in real terms before dividing by hours. It is an explicit top-code sensitivity, separate from the baseline prepared earnings treatment.

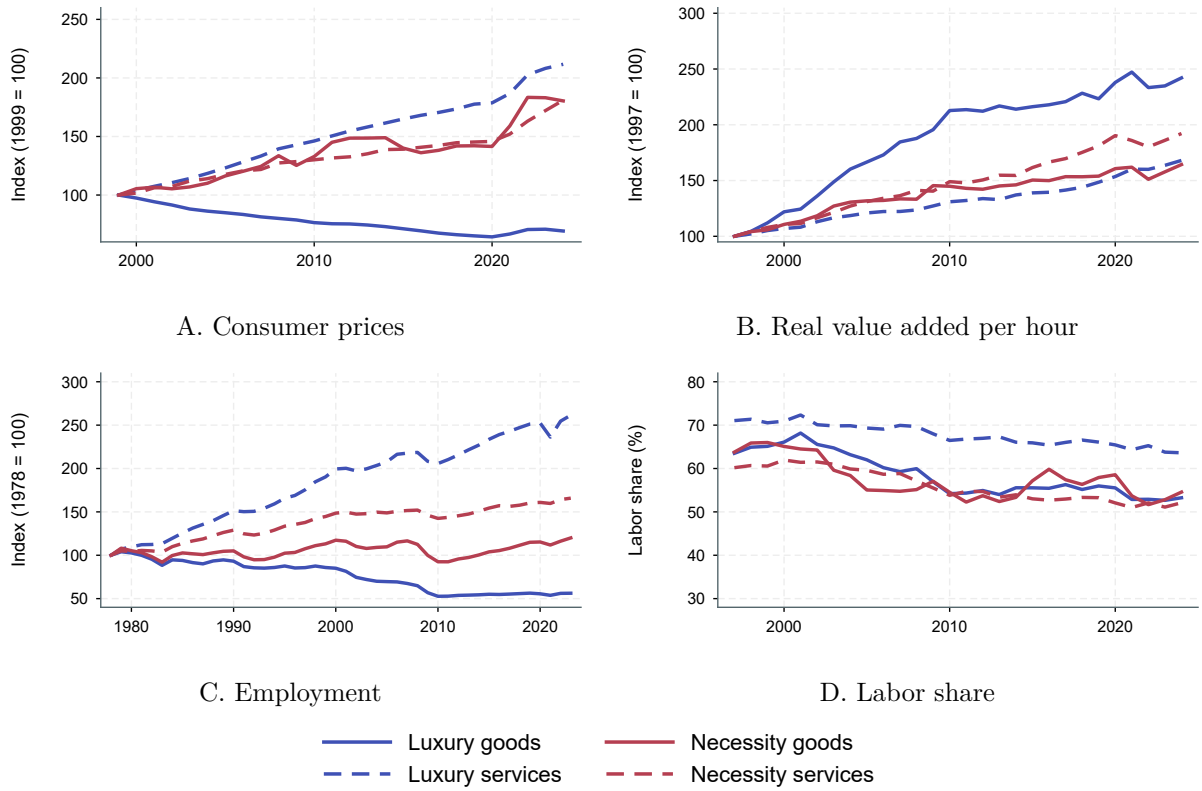
Direct reweighting to the 2003 joint age-band, sex, and education distribution provides an additional check where all required cells have support. Some cells in the higher-college-share industry group are empty in particular years, so I use regression standardization for that comparison rather than filling absent cells. All methods keep the original exposure weights and observed wage sample.

B Additional empirical evidence and robustness

B.1 Robustness to classification choices

The headline contrast between the two luxury sectors survives removing the full-coverage housing and healthcare assignments (Figure 18). LS prices and employment still rise fastest, while LG

Figure 18: Headline evidence excluding housing and healthcare



Notes: As in Figure 1, excluding housing and healthcare; the Engel specification and mapping method are unchanged.

prices fall and LG productivity grows fastest. Domestic LG employment continues to contract. LS productivity growth is slightly higher than in NG, but remains well below LG growth.

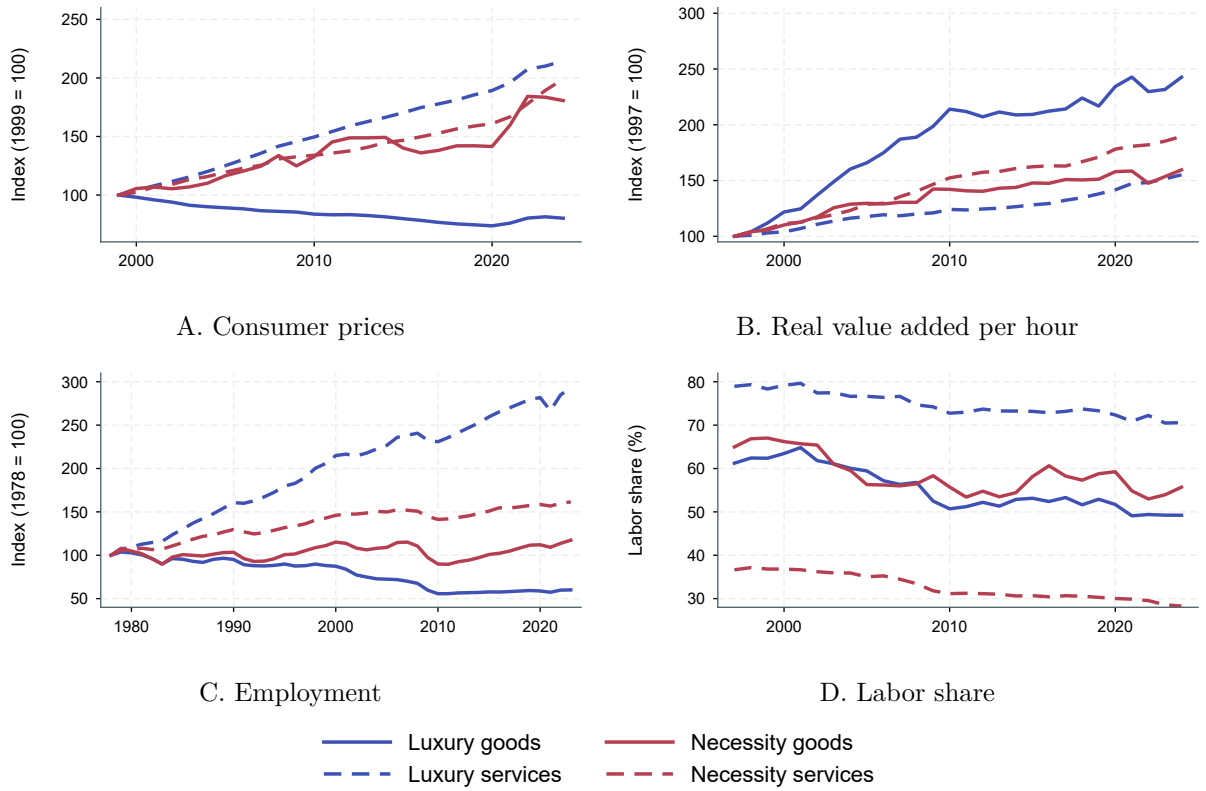
LS also retains the highest labor share (Panel D). In 2024, labor compensation accounts for 63.6% of LS value added, compared with 52.1–54.6% in the other three sectors. Relative to the baseline, the NS labor share rises from 30.8% to 52.1%, narrowing its gap with LS.

Engel Specification II yields the same ordering of price, productivity, and employment growth across the luxury sectors (Figure 19). This comparison retains the full-coverage treatment and changes the Engel specification. The contrast therefore survives both alternative category coverage and alternative estimated coefficients.

B.2 Alternative price and productivity measures

The contrast between luxury goods and luxury services survives alternative price and productivity measures (Figure 20). Panels A and B show slower price growth for luxury goods in both BEA PCE consumer-price indexes and KLEMS value-added deflators. Over 1999–2024, PCE prices for luxury goods fall while their value-added prices rise; the common pattern is slower price growth relative to luxury services. Panels C and D show faster productivity growth in luxury goods using integrated TFP and real value added per quality-adjusted labor input. The

Figure 19: Headline evidence under Engel Specification II



Notes: Figure 1 repeated using Engel Specification II. Housing and healthcare treatment, and all outcome definitions, are unchanged.

former accounts for growth in capital, labor, and intermediate inputs; the latter adjusts the labor denominator for workforce composition.

B.3 Reallocation and import exposure

Luxury services also expand faster than the other sectors in nominal value added and firm–industry counts (Figure 21). These measures complement the employment comparison. Nominal value added reflects changes in both quantities and prices. Firm counts are summed across four-digit industries, so a firm with establishments in several industries can be counted more than once, including within the same sector. Luxury-goods firm–industry counts decline over the period, while their nominal value added rises.

Final-use import exposure is highest in luxury goods, lower in necessity goods, and small in services (Figure 22). Hence, growing expenditure on goods need not translate into domestic production. These import shares are allocated using the make-table product-technology bridge rather than observed for each household. The main domestic-VA demand object removes both final-use and intermediate import content.

B.4 Luxury reallocation within goods and services

The luxury share of employment and nominal value added rises through changes within goods and services as well as the shift between them. Let G_t denote the goods share of either outcome

Figure 20: Alternative price and productivity measures



A. BEA PCE consumer prices (1960 = 100)

B. KLEMS value-added prices (1997 = 100)

C. KLEMS integrated TFP (1997 = 100)

D. VA per quality-adjusted labor (1997 = 100)

Notes: Panel A: BEA PCE Fisher price indexes, 1960–2024, adjusted for out-of-pocket healthcare. Panels B–D: KLEMS, 1997–2024. Integrated TFP uses Domar weights; quality-adjusted labor uses compensation weights. Allocation uses nominal value added, except gross output in Panel C. Indexes equal 100 in the indicated base year.

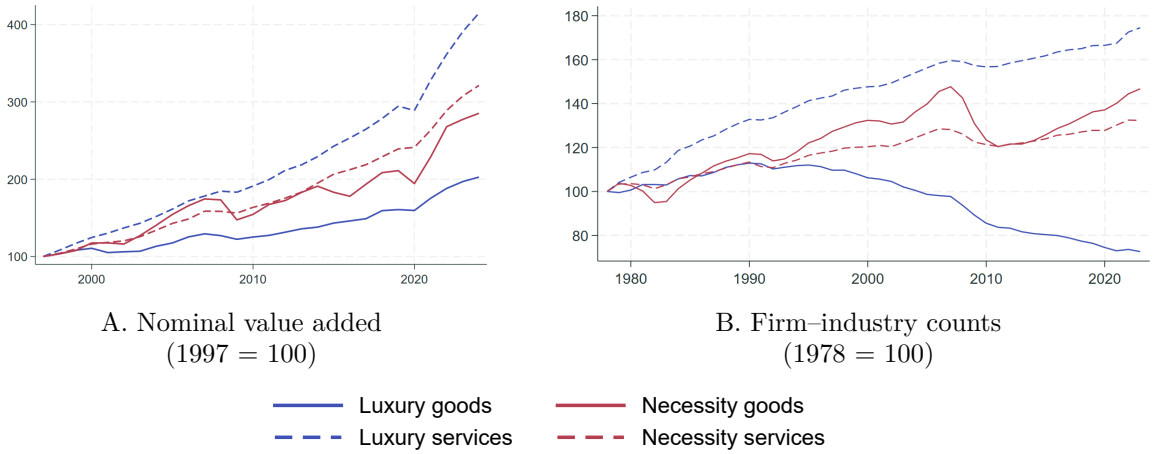
and let L_t^G and L_t^S denote its luxury shares within goods and services. The aggregate luxury share is

$$L_t = G_t L_t^G + (1 - G_t) L_t^S. \tag{30}$$

Figure 23 compares the observed share with two paths. Holding the initial within-group luxury shares fixed gives $G_t L_0^G + (1 - G_t) L_0^S$, the path implied by the goods-to-services shift alone. Holding the initial goods share fixed instead gives $G_0 L_t^G + (1 - G_0) L_t^S$, isolating changes in luxury shares within the two groups. Relative to the initial share, the observed change equals the sum of these two path changes plus the interaction $(G_t - G_0)[(L_t^G - L_0^G) - (L_t^S - L_0^S)]$. The separate comparison paths are therefore not an additive decomposition without that term.

Holding the initial luxury shares within each group fixed understates the observed increase in the luxury share of both employment and nominal value added, particularly the latter. These comparisons concern production-side activity; they do not establish an increase in the aggregate luxury share of CEX expenditure.

Figure 21: Reallocation in nominal value added and firm activity

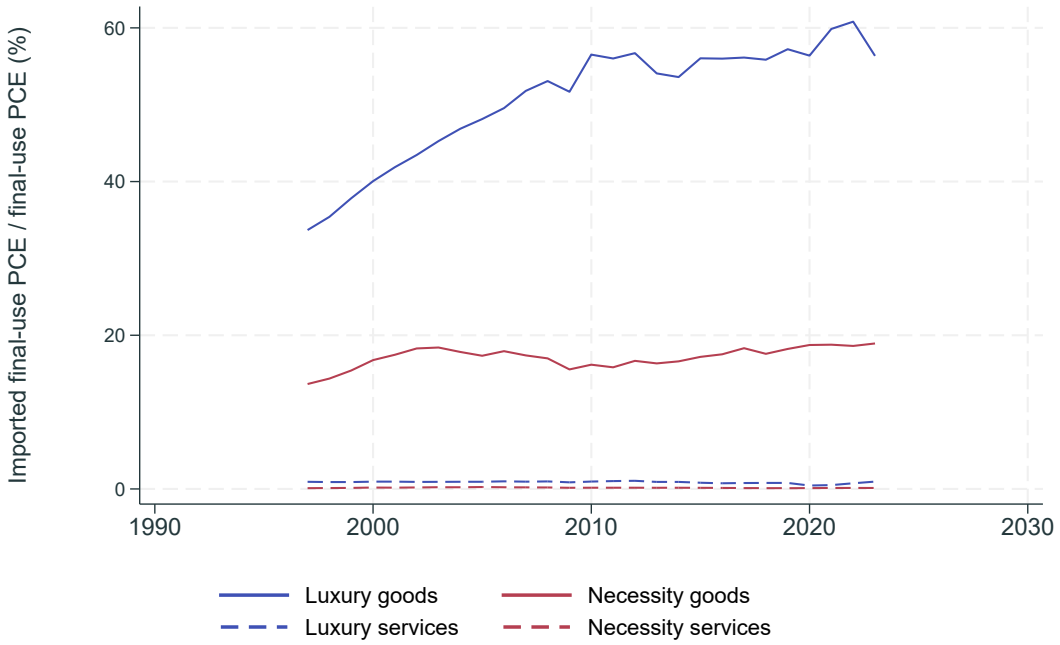


A. Nominal value added (1997 = 100)

B. Firm–industry counts (1978 = 100)

Notes: BEA industry accounts (value added, 1997–2024) and BDS (firms, 1978–2023); first year = 100. Classified activity only. Firm counts sum across industries, so multi-industry firms can appear more than once.

Figure 22: Import exposure of household expenditure



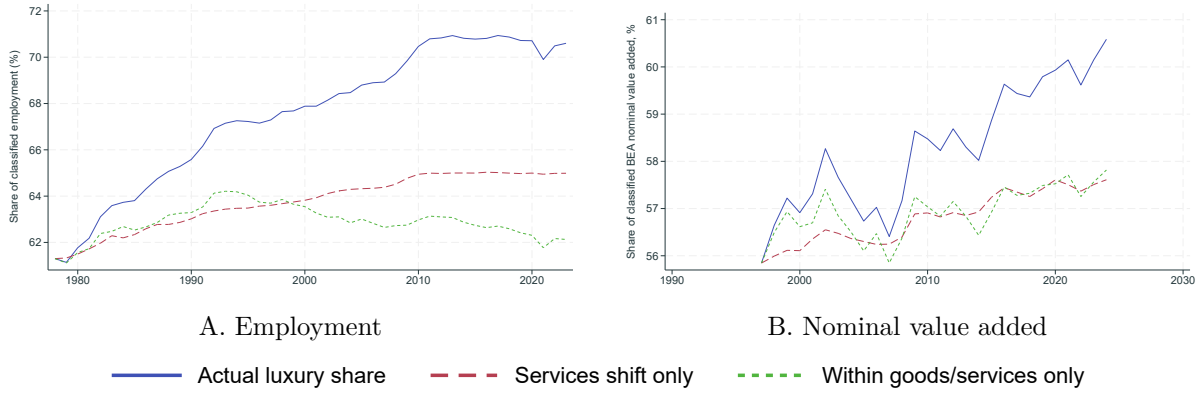
Notes: BEA PCE and input-output tables, 1997–2023; baseline classification and 2017 product-technology bridge. Final-use import exposure is imported final-use PCE divided by total final-use PCE, in current dollars, expressed as a percentage. Levels are summed within each sector before taking the ratio. Import allocations use the BEA import-ratio assumptions; the plotted measure excludes imported intermediate inputs.

B.5 Factor intensity, tasks, and worker comparisons

Luxury services have the lowest capital share of the four sectors, complementing the labor-share comparison in Figure 1, Panel D. These factor shares describe how value added is divided between labor and capital.

Luxury services have the lowest routine-task share and the highest share of highly interpersonal employment (Figure 25). Necessity services also rely more on interpersonal work than either

Figure 23: Luxury reallocation beyond the goods-to-services shift



Notes: BDS employment and BEA nominal value added. Comparison paths fix initial goods/services composition or initial luxury shares within each branch. Separate changes are not additive because of their interaction.

Figure 24: Capital shares of value added



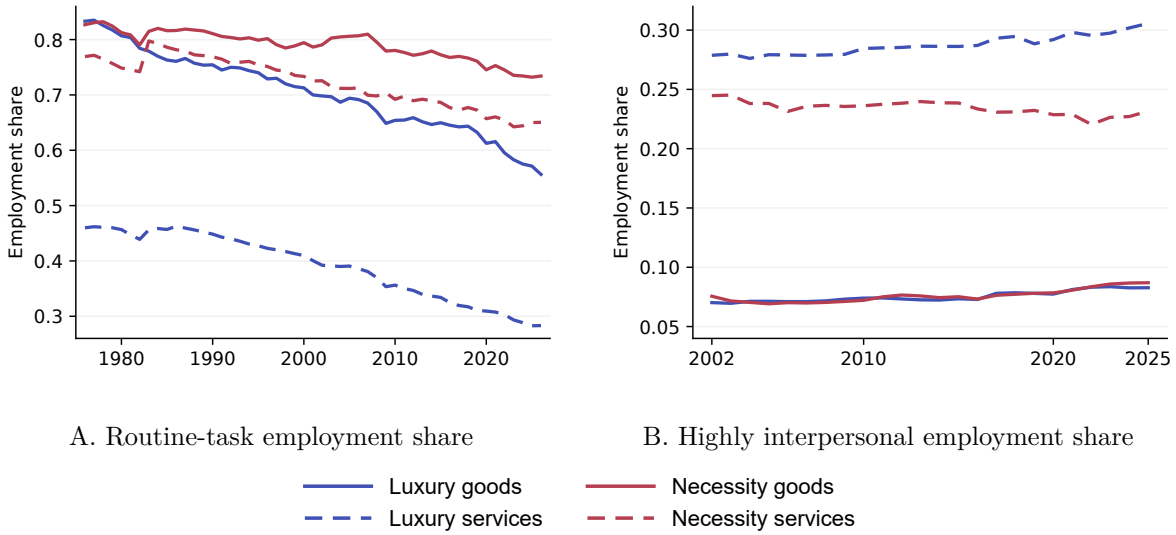
Notes: KLEMS, 1997–2024. Capital compensation divided by nominal value added, using baseline fractional industry allocation.

goods sector. These patterns help characterize service production, but do not measure the scope for future automation.

LS exposure declines across earnings deciles within both education groups (Figure 26). The rising college share alone would predict the opposite gradient: Holding each education group's LS share fixed makes aggregate LS exposure rise with earnings.

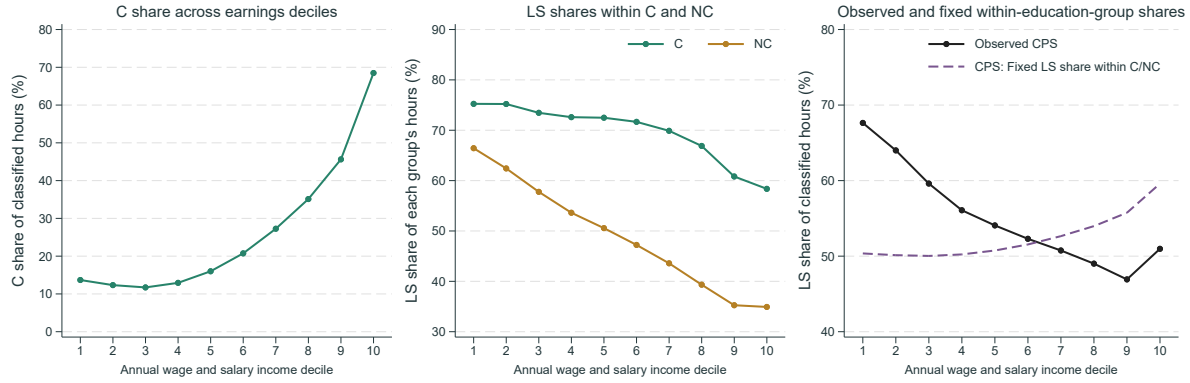
The hourly-wage comparison in Figure 27 addresses the role of hours in annual earnings. LS exposure falls across hourly-wage deciles within both education groups, while the college share rises. Fixing each education group's LS exposure at its pooled value again predicts a rising aggregate gradient. The within-group pattern therefore reflects exposure among lower-paid workers, not only differences in annual hours. These individual wage ranks differ from the fixed occupation groups ranked by initial pay in Figure 3.

Figure 25: Task content across the four sectors



Notes: CPS Basic Monthly, 1976–2026, private wage/salary workers (routine cognitive and routine manual tasks); BLS OEWS and O*NET, 2002–2025 (interpersonal tasks). Fractional four-sector allocation under the baseline classification. Each panel uses its source's employment weights and occupation-to-industry mapping.

Figure 26: Luxury-service worker exposure within education groups



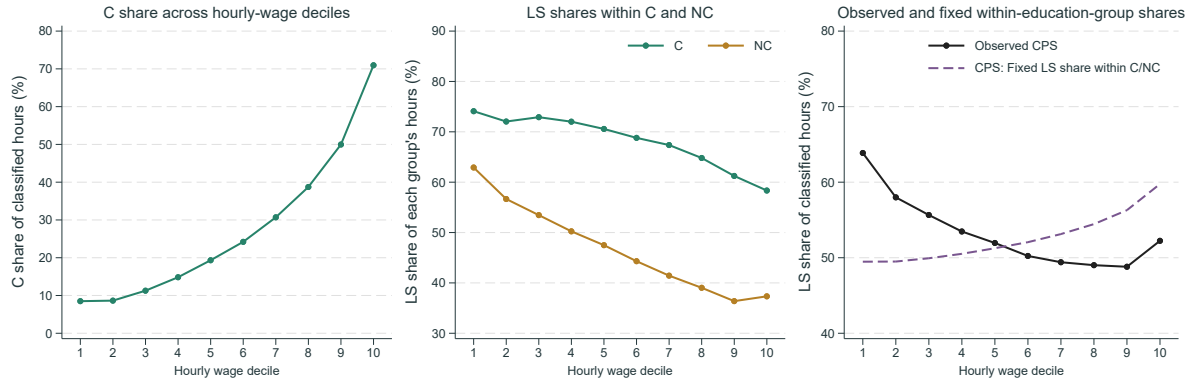
Notes: CPS ASEC, 1980–2024. Person-weighted annual earnings deciles; shares use classified hours. Dashed line: Each decile's observed education mix with pooled within-education LS exposure held fixed.

The employment and wage patterns also appear within fixed groups of industries defined by their initial college employment shares (Table 12). Over 2003–2024, non-college employment falls least in luxury services among industries with low initial college shares and rises most among industries with high initial college shares. Composition-adjusted hourly wages of non-college workers also grow most in luxury services within both industry groups. Appendix A.5 gives the fixed industry split and composition adjustment.

B.6 Consumption within age and education groups

Figure 28 asks whether the rise in LS spending occurs within broad demographic groups. From the 2000s to the 2010s, the LS share rises for households whose reference person is aged 60 or younger and for older households. It also rises within all five education groups. Changes in the relative sizes of these groups cannot be the entire explanation for that increase. Earlier trends

Figure 27: Luxury-service exposure across hourly-wage deciles



Notes: CPS ORG, 1982–2024; private employees aged 16–64. National hourly-wage deciles are defined each year across all workers using earnings weights, rather than separately by education; share weights also include usual hours. Dashed line: Observed education mix with pooled within-education LS exposure fixed. Denominators include only classified hours.

Table 12: Non-college workers: Employment and adjusted wage growth, 2003–2024

Industry group	Sector	Employment	Adjusted wage
Lower college share	NG	-3.53	0.50
	NS	-4.34	-0.22
	LG	-36.03	0.47
	LS	-0.87	4.44
Higher college share	NS	-15.61	4.48
	LG	-54.91	-13.75
	LS	7.49	5.59

Notes: CPS Basic Monthly (employment) and ORG (wages), 2003–2024; percentage changes. Industry groups use 2003–2005 college shares, with a 50% cutoff. Arithmetic wages hold 2003 age, sex, and non-college education composition fixed and apply a common real earnings ceiling (Appendix A.5). NG has no exposure in the higher-college-share group.

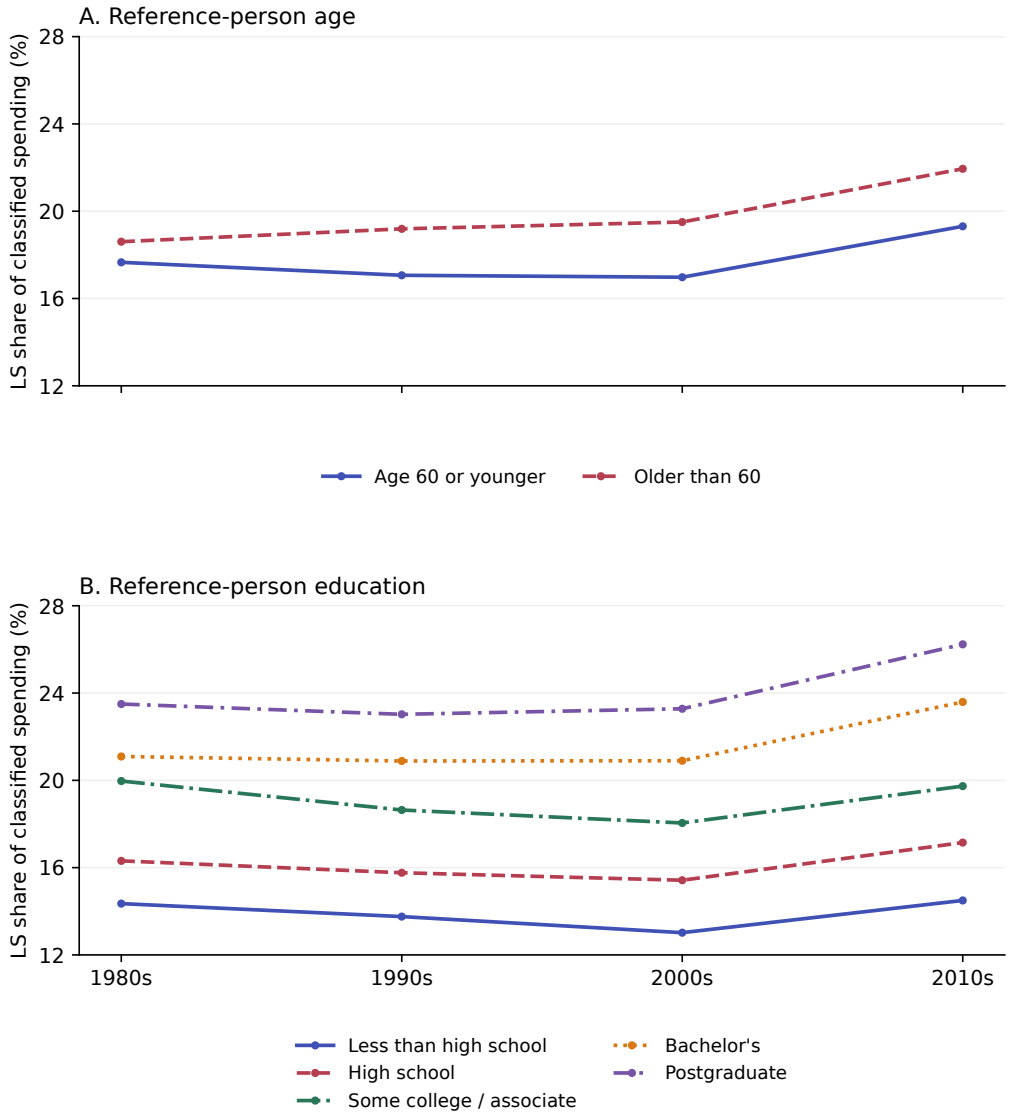
are mixed, however, and changes in income or household composition within these broad groups remain possible. These repeated cross-sections do not follow the same households over time.

The demand account offers a complementary comparison: It separates the contribution of changes in demographic and regional composition while retaining fixed estimated taste coefficients (Section 3 and Figure 34). Composition contributes to the predicted LS increase, but income and relative-price responses account for most of it. Neither exercise identifies a causal effect of age or education. Moreover, Figure 28 measures final consumption expenditure, whereas the demand account measures domestic consumption value added.

B.7 Growth accounting: Derivation and robustness

Constructing the four sectors. The baseline aggregates activity to the four sectors before computing the observed and fixed-weight accounts. Let $a_{js,t}$ be the fraction of native industry j assigned to sector s, and let $\mathcal{S} = \{\text{NG}, \text{NS}, \text{LG}, \text{LS}\}$. With superscript n denoting native

Figure 28: Luxury-service spending within age and education groups



Notes: CEX, 1980–2019; survey-weighted shares of classified nominal interview spending. Age panel: Equal-weighted annual means within decades. Education panel: Pooled interviews. Reference-person characteristics; missing education is excluded. Postgraduate means more than four college years.

quantities, sectoral nominal value added and hours are

$$V_{st} = \sum_j a_{js,t} V_{jt}^n, \qquad H_{st} = \sum_j a_{js,t} H_{jt}^n. \tag{31}$$

The classified fraction of industry j is $c_{jt} = \sum_{s \in \mathcal{S}} a_{js,t}$. One wholly unclassified private industry is excluded, leaving 60 native industries contributing to the four sectors. Pre-1997 allocations use the 1997 exposures. Historical KLEMS supplies growth through the 1997 link; current KLEMS supplies subsequent growth, with its industry hours indexes anchored to historical hours levels in 1997.

Real value added is chained within each sector using midpoint allocated nominal-VA shares. Define $q_{js,t} = a_{js,t}V^n_{jt}/V_{st}$ and $\bar{q}_{js,t} = (q_{js,t-1} + q_{js,t})/2$. Then

$$\Delta \log Y_{st} = \sum_j \bar{q}_{js,t} \left(\Delta \log Y^n_{jt} + \Delta \log a_{js,t} \right). \tag{32}$$

The sum covers industry-sector cells with positive allocation at both endpoints; permanently zero cells carry zero weight. The inputs contain no transitions between zero and positive industry-sector allocations. Thus, changing classification exposure scales allocated real output as well as nominal value added and hours. The account does not sum nonadditive chained-dollar levels. Sectoral productivity growth is $g^{\text{LP}}_{st} = \Delta \log Y_{st} - \Delta \log H_{st}$. Both the observed and fixed accounts retain this path, including its within-sector composition changes.

The gap and its sectoral contributions. Let $\omega^V_{st} = V_{st}/\sum_r V_{rt}$ and $\bar{s}^V_{st} = (\omega^V_{s,t-1} + \omega^V_{st})/2$. The fixed weight $\bar{s}^V_{s0}$ is the 1980–1981 midpoint. For hours, define $H_t = \sum_s H_{st}$, $u_{st} = H_{st}/H_t$, and the logarithmic mean $\mathcal{L}(a,b) = (a-b)/(\log a - \log b)$, with $\mathcal{L}(a,a) = a$. The weights

$$h_{st} = \frac{\mathcal{L}(u_{st}, u_{s,t-1})}{\sum_r \mathcal{L}(u_{rt}, u_{r,t-1})} \tag{33}$$

give exact aggregate hours growth: $\sum_s h_{st} \Delta \log H_{st} = \Delta \log H_t$. Subtracting equation (1) from equation (2) gives

$$\begin{aligned} g^{\text{fixed}}_t - g^{\text{obs}}_t = &\sum_s (\bar{s}^V_{s0} - \bar{s}^V_{st}) g^{\text{LP}}_{st} \\ &+ \sum_s (h_{st} - \bar{s}^V_{st}) \Delta \log u_{st}. \end{aligned} \tag{34}$$

The second term can equivalently use $\Delta \log H_{st}$ because $\sum_s (h_{st} - \bar{s}^V_{st}) = 0$. Centering hours growth on aggregate hours growth therefore leaves the total gap unchanged, while making each sector's hours contribution depend on changes in its hours share. Hence, unchanged hours shares imply zero hours-reallocation contributions, even when aggregate hours grow.

To allocate the value-added-weight component across sectors, I express sectoral productivity growth relative to a common reference. This reference averages productivity growth evaluated at initial and current value-added weights, treating the two sets of weights symmetrically:

$$b_t = \frac{1}{2} \left[\sum_s \bar{s}^V_{s0} g^{\text{LP}}_{st} + \sum_s \bar{s}^V_{st} g^{\text{LP}}_{st} \right]. \tag{35}$$

The reference differs from the average of observed and fixed aggregate growth because observed growth also contains hours reallocation. The sectoral contributions are

$$D^V_{s,t} = (\bar{s}^V_{s0} - \bar{s}^V_{st})(g^{\text{LP}}_{st} - b_t), \tag{36}$$

$$D^H_{s,t} = (h_{st} - \bar{s}^V_{st})[\Delta \log H_{st} - \Delta \log H_t] = (h_{st} - \bar{s}^V_{st}) \Delta \log u_{st}, \tag{37}$$

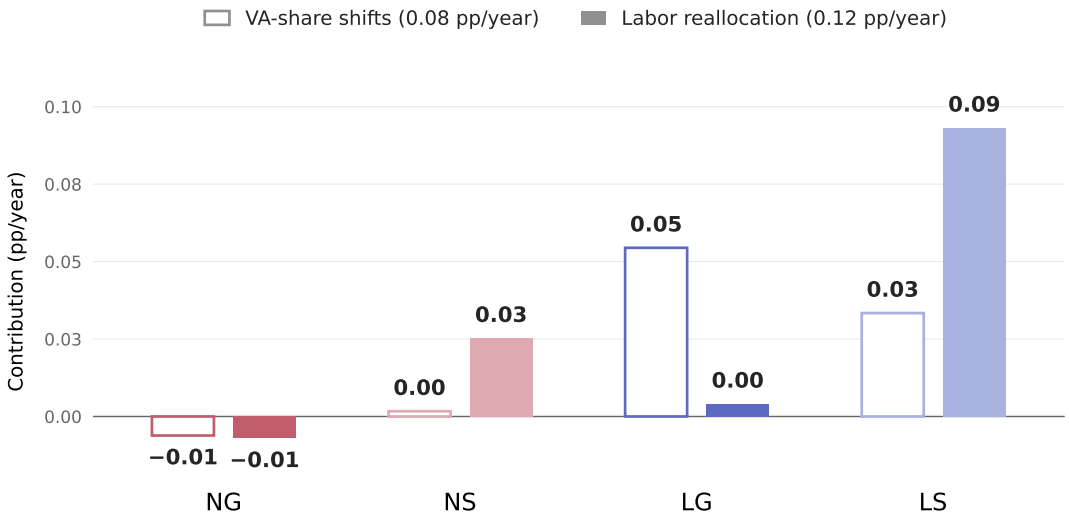
$$D_{s,t} = D^V_{s,t} + D^H_{s,t}, \qquad \sum_{s \in \mathcal{S}} D_{s,t} = g^{\text{fixed}}_t - g^{\text{obs}}_t. \tag{38}$$

Both value-added-weight vectors sum to one, so subtracting b_t leaves the aggregate value-added-weight component unchanged. The reference affects its allocation across sectors, not the aggregate gap. An expansion relative to the initial value-added share contributes positively when productivity growth is below b_t; a contraction contributes positively when growth is above b_t. The hours term instead pairs changes in hours shares with the difference between exact hours-aggregation and midpoint nominal-value-added weights. In continuous time, writing $z_s = V_s/H_s$ and $Z = V/H$, its counterpart is $(1 - z_s/Z)\,\mathrm{d}u_s$. Thus, expansion of a sector with below-average nominal value added per hour, or contraction of one with above-average nominal value added per hour, contributes positively through this margin.

I multiply annual log-point contributions by 100 and average across links to report percentage points per year. Contributions refer to fixed minus observed growth. This attribution specifies an additive accounting allocation; it does not identify causal sector effects or isolate the income-driven part of the luxury shift.

Figure 29 separates the four-sector gap into value-added-weight shifts and hours reallocation. Luxury services contribute 0.13 percentage points per year, comprising 0.03 from the value-added term and 0.09 from the hours term. Luxury goods contribute mainly through value-added-weight shifts, while the positive necessity-service contribution comes mainly from hours reallocation. Necessity goods have small negative contributions through both margins.

Figure 29: Components of the four-sector productivity gap



Notes: Merged KLEMS, 1980–2024; classified private activity under the baseline classification. Average annual contributions to fixed-minus-observed growth, in percentage points per year. The benchmark fixes 1980–1981 midpoint value-added weights. Hollow bars show value-added-share shifts using relative attribution; filled bars show labor reallocation using centered hours-share growth. Positive values widen the gap.

Table 13 compares 1981–2004 with 2005–2024, using the original 1980–1981 midpoint value-added weights in both periods. Observed productivity growth falls from 2.08% to 1.54% per year, while growth in the fixed-weight benchmark falls from 2.35% to 1.65%. Hence, observed growth slows by less than benchmark growth, and the gap narrows from 0.27 to 0.11 percentage

points per year. The luxury-service contribution remains positive on average in both periods, but falls from 0.17 to 0.07 percentage points per year.

Table 13: Four-sector productivity growth by period

Period	Observed	Fixed	Gap	LS
1981–2004	2.08	2.35	0.27	0.17
2005–2024	1.54	1.65	0.11	0.07

Notes: Merged KLEMS; classified private activity. Both periods use fixed 1980–1981 midpoint value-added weights. Observed and Fixed report mean annual log growth in percent. Gap is fixed minus observed growth; LS is the luxury-service contribution to the gap, both in percentage points per year. Periods refer to the ending years of annual changes.

Detailed-industry comparison. The comparison account fixes weights across 60 KLEMS industries. It first removes unclassified activity within each industry: $V_{jt}^N = c_{jt}V_{jt}^n$, $H_{jt}^N = c_{jt}H_{jt}^n$, and $\Delta \log Y_{jt}^N = \Delta \log Y_{jt}^n + \Delta \log c_{jt}$. Proportional allocation leaves each industry's productivity growth unchanged:

$$g_{jt}^{\mathrm{LP},N} = \Delta \log Y_{jt}^N - \Delta \log H_{jt}^N = \Delta \log Y_{jt}^n - \Delta \log H_{jt}^n. \tag{39}$$

The observed and fixed formulas have the same form as equations (1)–(2), replacing sectors by industries. Denote the corresponding industry value-added weights, exact hours weights, hours shares, and symmetric productivity-growth reference by $\bar{s}_{jt}^{V,N}$, h_{jt}^N, u_{jt}^N, and b_t^N.

Industry contributions are assigned to the four sectors using normalized midpoint exposures,

$$e_{js,t} = \frac{a_{js,t-1} + a_{js,t}}{c_{j,t-1} + c_{jt}}, \qquad \sum_{s \in \mathcal{S}} e_{js,t} = 1. \tag{40}$$

This normalizes the midpoint raw exposures by the midpoint classified fraction; it need not equal the average of separately normalized endpoint exposures. The native-industry sector contributions are

$$D_{s,t}^{V,N} = \sum_j e_{js,t}(\bar{s}_{j0}^{V,N} - \bar{s}_{jt}^{V,N})(g_{jt}^{\mathrm{LP},N} - b_t^N), \tag{41}$$

$$D_{s,t}^{H,N} = \sum_j e_{js,t}(h_{jt}^N - \bar{s}_{jt}^{V,N})\Delta \log u_{jt}^N. \tag{42}$$

They sum to the native-industry gap in every annual link. The hours term is again centered on actual aggregate hours growth.

Table 14 compares the two accounts. The detailed-industry gap is 0.22 percentage points per year, compared with 0.20 in the four-sector account. The luxury-service contribution is 0.13 in both, but its industry-level hours term is larger (0.12 rather than 0.09) and its value-added term smaller (0.01 rather than 0.03). The other sectors' allocations change more: The detailed-industry account assigns a larger contribution to necessity goods and a small negative contribution to luxury goods. The full-period gap and luxury-service contribution are therefore similar across aggregation levels, although their timing and the allocation to other sectors differ.

Table 14: Four-sector and detailed-industry growth accounting

A. Aggregate annual growth, 1980–2024

Accounting units	Observed	Fixed	Gap
Four sectors	1.83	2.03	0.20
60 industries	1.83	2.05	0.22

B. Sectoral contributions to the gap

	Four sectors			60 industries		
Sector	VA shifts	Hours	Total	VA shifts	Hours	Total
NG	-0.01	-0.01	-0.01	0.05	0.01	0.05
NS	0.00	0.03	0.03	0.03	0.02	0.05
LG	0.05	0.00	0.06	0.00	-0.01	-0.01
LS	**0.03**	**0.09**	**0.13**	**0.01**	**0.12**	**0.13**
Total	**0.08**	**0.12**	**0.20**	**0.08**	**0.14**	**0.22**

Notes: KLEMS, 1980–2024; classified private activity. Weights are fixed after aggregation in the four-sector account and before aggregation in the detailed-industry account. Both use 1980–1981 midpoint weights. Growth: Annual log percent. Gaps and contributions: Percentage points per year.

The difference between the two gaps equals the difference in fixed-benchmark growth minus the difference in observed growth. Observed growth is 1.83% in both accounts at the reported precision. Most of the gap difference comes from the fixed benchmark: Industry weights are fixed in the detailed-industry account, whereas within-sector composition continues to change along the observed sectoral productivity paths in the four-sector account. Chaining midpoint weights in two stages also produces a small difference in observed real-value-added growth. Hence, subtracting the two gaps does not give an exact within-sector reallocation term.

Alternative labor-productivity accounts. Starting the four-sector account in 1997 and resetting the fixed weights to the 1997–1998 midpoint reduces the annual gap from 0.20 to 0.14 percentage points, but luxury services remain the largest contributor at 0.08 (Table 15). Replacing total hours by quality-adjusted labor services in the four-sector account over 1980–2024 gives observed growth of 1.40% and benchmark growth of 1.59% per year. Luxury services contribute 0.13 percentage points to the 0.19-point gap, again the largest sectoral contribution.

The labor-services comparison uses the merged KLEMS industry labor-input index, L^n_{jt}, which adjusts hours for workforce composition. Let C^n_{jt} denote industry labor compensation and $C_{st} = \sum_j a_{js,t} C^n_{jt}$ its allocated sector total. Define within-sector compensation shares $\ell_{js,t} = a_{js,t} C^n_{jt} / C_{st}$ and sector shares $\ell_{st} = C_{st} / \sum_r C_{rt}$, with bars denoting adjacent-year midpoints. Sectoral and aggregate labor-index growth are

$$\begin{aligned} \Delta \log L_{st} &= \sum_j \bar{\ell}_{js,t} \left(\Delta \log L^n_{jt} + \Delta \log a_{js,t} \right), \\ \Delta \log L_t &= \sum_s \bar{\ell}_{st} \Delta \log L_{st}. \end{aligned} \tag{43}$$

Classified exposures allocate nominal value added and compensation before weights are formed; changes in exposure enter both real-output and labor-services growth. Zero allocation cells

follow the same rule as in equation (32), and signed industry compensation is retained. The account keeps the baseline's sectoral real-output paths and 1980–1981 value-added weights, replacing hours by labor services. Its labor contribution is $(\bar{\ell}_{st} - \bar{s}^{V}_{st})[\Delta \log L_{st} - \Delta \log L_t]$: Midpoint compensation shares replace exact-hours weights, and sectoral labor growth is centered on aggregate labor-index growth. The value-added-weight component uses the same symmetric reference, applied to productivity per unit of labor services. The denominator is therefore a chained index of quality-adjusted labor services.

Table 15: Alternative labor-productivity accounts

Account	Observed	Fixed	Gap	NG	NS	LG	LS
Alternative starting year: 1997–2024	1.90	2.05	0.14	-0.01	0.04	0.03	0.08
VA/labor services, 1980–2024	1.40	1.59	0.19	-0.01	0.04	0.03	0.13

Notes: Classified private activity. Row 1: Current KLEMS, total hours, 1997–2024. Row 2: Merged KLEMS, quality-adjusted labor, 1980–2024. Benchmarks use first-link midpoint weights. Growth is in annual log percent; gaps and contributions in percentage points per year.

C Demand estimation and accounting

C.1 Expenditure shares and Engel elasticities

Let S_{ij} be household i's expenditure share in outer nest j and let q_{ijk} be its conditional share in good or service k. Shephard's lemma applied to equations (3)–(4) gives

$$q_{ijk} = \frac{\beta_{ijk} u_i^{\varepsilon_{jk}-1} p_{jk}^{1-\sigma_j}}{\sum_m \beta_{ijm} u_i^{\varepsilon_{jm}-1} p_{jm}^{1-\sigma_j}}, \tag{44}$$

$$S_{ij} = \frac{\beta_{ij} u_i^{\varepsilon_j - 1} P_{ij}^{1-\sigma_0}}{\sum_h \beta_{ih} u_i^{\varepsilon_h - 1} P_{ih}^{1-\sigma_0}}, \qquad s_{ijk} = S_{ij} q_{ijk}. \tag{45}$$

Define $a_{ij} = \sum_k q_{ijk}(\varepsilon_{jk} - 1)/(1-\sigma_j)$ and $B_{ij} = \varepsilon_j - 1 + (1-\sigma_0)a_{ij}$. The elasticity of expenditure with respect to utility is

$$\zeta_i = \frac{\partial \log \mathcal{E}_i}{\partial \log u_i} = 1 + \sum_j S_{ij} \left(\frac{\varepsilon_j - 1}{1 - \sigma_0} + a_{ij} \right). \tag{46}$$

At fixed prices, a sector's expenditure elasticity is

$$\eta_{ijk} = 1 + \frac{B_{ij} - \sum_h S_{ih} B_{ih} + \varepsilon_{jk} - 1 - (1-\sigma_j) a_{ij}}{\zeta_i}. \tag{47}$$

The common outer share cancels from $s_{i,\mathrm{LS}}/s_{i,\mathrm{LG}}$, giving equation (6).

The industry classification is fixed using category-level estimates, whereas the model sector's local Engel elasticity varies with expenditure shares. Aggregation and the household's position on its Engel curve therefore matter for the latter.

C.2 Estimating equations and instruments

The estimated equations use within-nest and between-nest expenditure ratios. Writing ℓ_i for normalized log utility and D_i for observed taste controls,

$$\log \frac{x_{i\mathrm{NS}}}{x_{i\mathrm{NG}}} = D_i'\gamma_N + (\varepsilon_{\mathrm{NS}} - 1)\ell_i + (1 - \sigma_N) \log \frac{p_{i\mathrm{NS}}}{p_{i\mathrm{NG}}} + \xi_{iN}, \tag{48}$$

$$\log \frac{x_{i\mathrm{LS}}}{x_{i\mathrm{LG}}} = D_i'\gamma_L + (\varepsilon_{\mathrm{LS}} - 1)\ell_i + (1 - \sigma_L) \log \frac{p_{i\mathrm{LS}}}{p_{i\mathrm{LG}}} + \xi_{iL}, \tag{49}$$

$$\log \frac{x_{iL}}{x_{iN}} = D_i'\gamma_0 + (\varepsilon_L - 1)\ell_i + (1 - \sigma_0) \log \frac{P^*_{iL}}{P^*_{iN}} + \xi_{i0}. \tag{50}$$

Here $x_{ij} = x_{ijG} + x_{ijS}$, $E_i = x_{iN} + x_{iL}$, and

$$P^*_{ij} = p_{ijG} \left(\frac{x_{ijG}}{x_{ij}} \right)^{-1/(1-\sigma_j)}, \qquad \ell_i = \log \frac{E_i}{P^*_{iN}} + \frac{\log(x_{iN}/E_i)}{1 - \sigma_0}. \tag{51}$$

Taste controls include a constant, indicators for the household head's age 38–50 and 51 or older, household size 3–4 and 5 or more, and two or more earners. Column (2) of Table 2 adds region fixed effects. Column (3) also includes calendar controls for 2000–2023. The displayed columns (1)–(3) correspond to saved specifications 1, 3, and 5.

Every equation uses the same instrument block: The included taste controls, four before-tax-income-class indicators (classes 2–5, with class 1 omitted), log adjusted real after-tax income, and all three other-region log price ratios NS/NG, LS/LG, and LG/NG. Income classes retain their historical and continuation conventions. The logged income measure combines the final-interview annual component, including non-income and retirement adjustments, with four quarterly adjustments for housing, selected expenses, and food stamps. Thus, it is an adjusted resource measure rather than independently observed cash income. Both income classes and this log-income measure are constant across a household's retained interviews. Other-region component CPIs are aggregated with contemporaneous own-region expenditure weights. All excluded instruments require orthogonality to residual taste shocks conditional on the controls; the income moments do not by themselves remove expenditure-reporting error or unobserved preference heterogeneity.

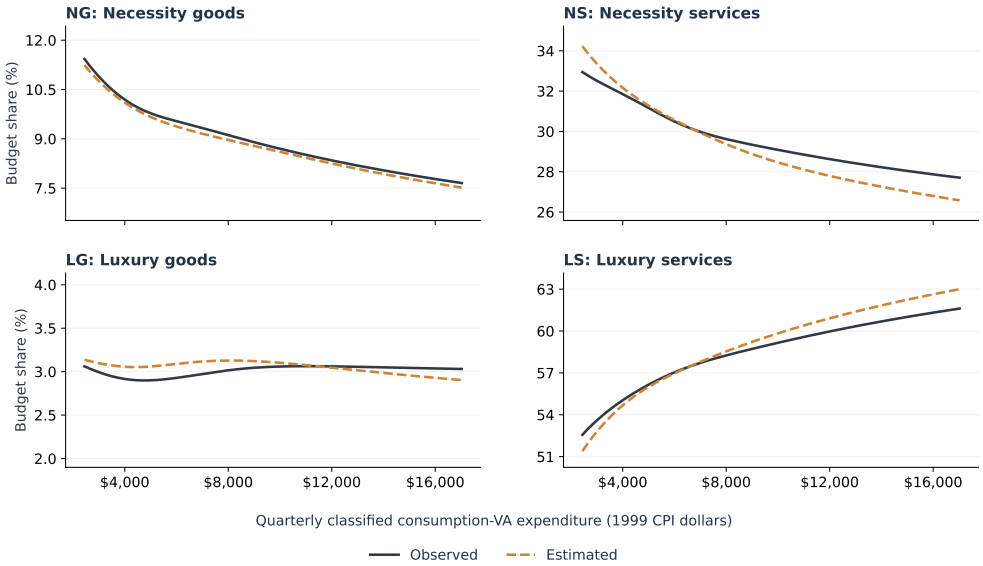
The equations are estimated jointly using survey-weighted GMM moments and household-clustered inference. Complete interview panels are formed before expenditure trimming, and household-specific consumer-price exposures are normalized consistently with the estimating equations. In fixed-taste historical predictions, calendar terms are held at their omitted category while demographic and regional characteristics can change. Column (3) is the baseline. Appendix C.6 restores its estimated calendar effects while retaining the same initial-share calibration and other coefficients.

C.3 Smooth Engel profiles in real dollars

Figure 30 complements the decile comparison in Figure 6 by placing observed and predicted budget shares on a common real-dollar expenditure axis. It uses the same preferred estimates,

eligible household observations, and saved predictions. I fit survey-weighted restricted cubic splines in log real expenditure separately within each year, using five knots and the same design for observed and predicted shares. Fixed survey-mass year weights then combine the profiles over 2000–2023.

Figure 30: Observed and predicted budget shares by real expenditure



Notes: CEX, 2000–2023; preferred domestic consumption-VA specification, within-year 2nd–98th percentile treatment. Predictions use observed household prices and characteristics and include the estimated calendar effects at each observation date. Expenditure is quarterly classified consumption value added in 1999 CPI dollars, using interview-quarter CPI deflation. Survey-weighted within-year five-knot restricted cubic splines use all eligible observations; the displayed range is the intersection of annual weighted 1st–99th percentile expenditure support. Fixed year weights combine observed and predicted profiles at common dollar levels.

The predictions use observed household prices and characteristics and include the estimated calendar effects at each observation date. These are descriptive smooth cross-sectional profiles, rather than structural Engel curves holding prices and characteristics fixed. Combining years on a common dollar axis makes the expenditure levels interpretable, but does not by itself establish that Engel curves are stable across periods.

C.4 Demand estimates including imports

Substitution remains below one at all three nests when I include imports in consumption value added. Table 16 pairs each domestic specification from Table 2 with its imports-included counterpart. Including imports raises the substitution estimates: In the preferred specification, σ_0 rises from 0.41 to 0.48, σ_N from 0.20 to 0.61, and σ_L from 0.24 to 0.58. The outer luxury and luxury-service nonhomotheticity parameters remain above one in all three specifications. Thus, the qualitative demand pattern survives including imports, although the magnitudes depend on the consumption measure. The estimation samples also differ because expenditure trimming is applied separately to each measure. I retain domestic consumption value added in the general-equilibrium model because its sectors produce domestic value added.

Table 16: Demand estimates: Domestic value added and value added including imports

Specification	(1)		(2)		(3)	
	Domestic	Including imports	Domestic	Including imports	Domestic	Including imports
σ_0	0.44	0.48	0.46	0.54	0.41	0.48
	(0.01)	(0.02)	(0.01)	(0.02)	(0.02)	(0.03)
σ_N	0.17	0.45	0.31	0.52	0.20	0.61
	(0.02)	(0.01)	(0.02)	(0.01)	(0.03)	(0.02)
σ_L	0.21	0.55	0.19	0.54	0.24	0.58
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
ε_L	1.42	1.62	1.56	1.92	1.41	1.86
	(0.02)	(0.04)	(0.03)	(0.07)	(0.03)	(0.10)
ε_{NS}	1.28	1.30	1.36	1.41	1.29	1.40
	(0.01)	(0.02)	(0.02)	(0.03)	(0.02)	(0.04)
ε_{LS}	1.81	1.52	1.94	1.69	1.81	1.66
	(0.02)	(0.03)	(0.03)	(0.04)	(0.03)	(0.06)
Region fixed effects	No	No	Yes	Yes	Yes	Yes
Calendar controls	No	No	No	No	Yes	Yes
Estimation weight	Survey	Survey	Survey	Survey	Survey	Survey
Interviews	269,518	270,193	269,518	270,193	269,518	270,193
Households	71,938	71,985	71,938	71,985	71,938	71,985

Notes: CEX, 2000–2023; survey-weighted nonlinear GMM. Specifications follow Table 2; expenditure trimming is applied separately, so samples differ. Domestic value added excludes final and intermediate imports. Household-clustered standard errors in parentheses.

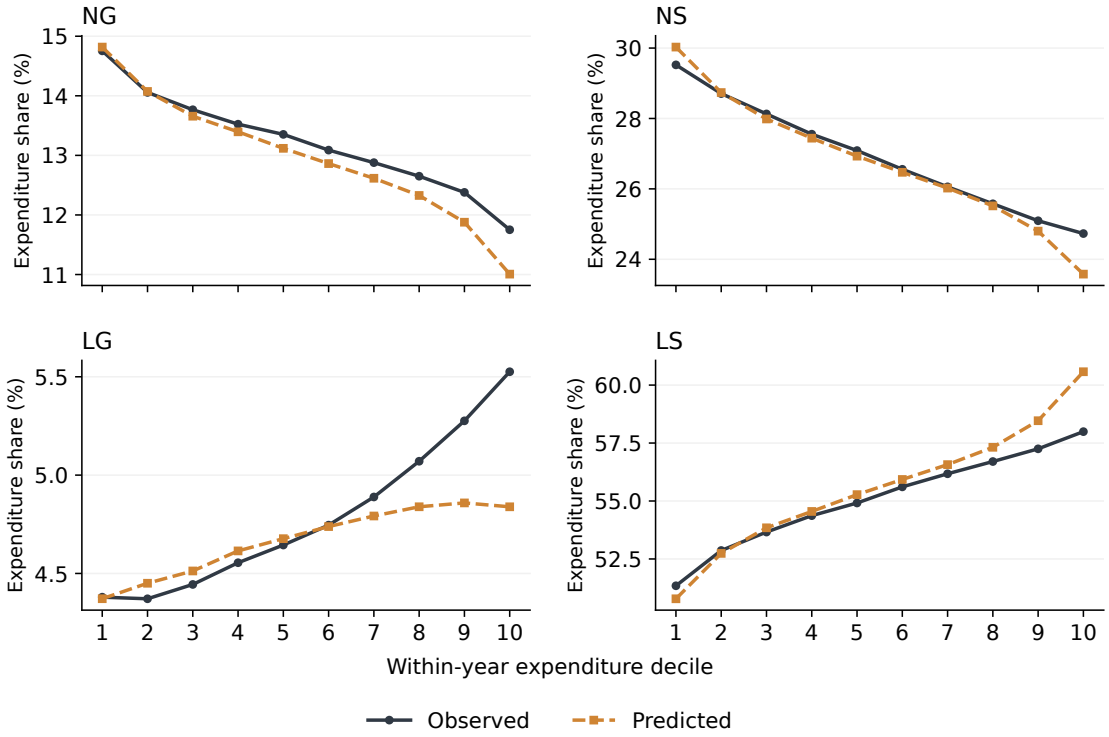
The imports-included estimates capture the declining necessity shares and rising luxury-service share across expenditure deciles (Figure 31). The fit is less close at the top of the expenditure distribution: In the highest decile, the predicted LS share exceeds the observed share by 2.58 percentage points, while the increase in the LG share is too small.

The demand system also reproduces most of the increase in the luxury-service share over time. Between 1999 and 2023, the observed share rises by 2.70 percentage points and the predicted share by 2.12 percentage points (Figure 32). It captures the decline in both necessity shares but understates the decline in luxury goods. These predictions use observed household budgets and common annual consumer-price indexes. Three taste intercepts match the initial sectoral shares, while demographic and regional coefficients remain fixed and calendar effects stay at their omitted reference; later shares therefore provide a check on the system's fit.

C.5 Aggregation and the decomposition of share changes

Figure 33 reports the price paths actually used in the empirical demand account. I form common annual prices as survey-weighted geometric means of normalized household estimation-price exposures on the matched accounting sample and rebase each sector to 1999. These are prices for the domestic consumption-VA object, constructed through the fixed input-output mapping. They differ from the headline D-CPI final-expenditure indexes and from the KLEMS producer prices used to assess the model. In particular, LG accounting prices rise over the sample, although less than prices in the other three sectors.

Figure 31: Observed and predicted Engel curves including imports



Notes: CEX, 2000–2023; survey-weighted consumption-VA shares including imports by annual real-expenditure decile. Predictions use the imports-included column (3) of Table 16, with observed household prices and characteristics and estimated calendar effects evaluated at each observation date.

For each adjacent-year comparison, I order both cross-sections by expenditure and form a common partition of their cumulative population distributions. Each interval inherits the budget and characteristics of its source observation at either endpoint. Splitting observation weights where necessary preserves both endpoint distributions and gives fixed population masses π_i along the link. The construction does not impose a household panel interpretation.

With expenditure weights $\nu_i = \pi_i x_i / \sum_h \pi_h x_h$ and fixed π_i,

$$d\nu_i = \nu_i \left(d \log x_i - \mathbb{E}_\nu[d \log x_i] \right). \tag{52}$$

Differentiating $\widehat{S}_s = \sum_i \nu_i s_{is}$ therefore gives

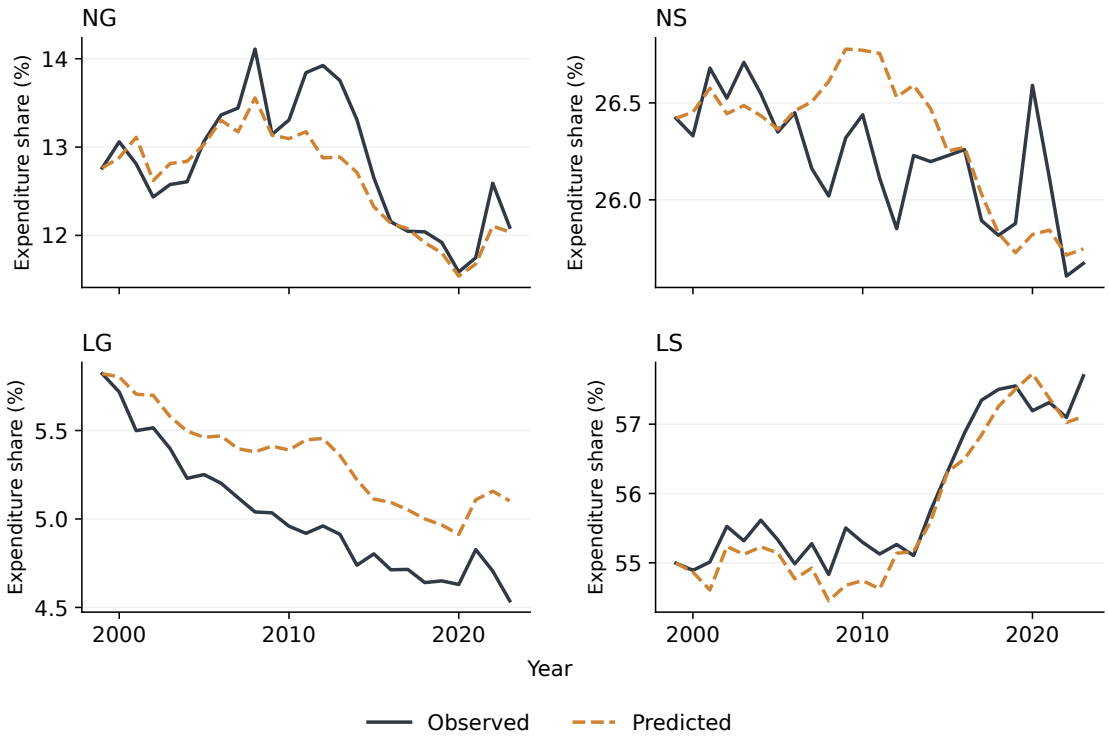
$$d\widehat{S}_s = \mathbb{E}_\nu[ds_{is}] + \mathrm{Cov}_\nu(s_{is}, d \log x_i). \tag{53}$$

The complete household-share differential is

$$ds_{is} = s_{is} B_{is} + s_{is}(\eta_{is} - 1) \left(d \log x_i - \sum_r s_{ir} d \log p_r \right) + T_{is}. \tag{54}$$

Here T_{is} is the fixed-budget share-level response to changing characteristics, with the estimated taste coefficients fixed. Substituting this expression into equation (53) gives the four-term iden-

Figure 32: Observed and predicted sectoral expenditure shares including imports



Notes: CEX, 1999–2023; consumption value added including imports, using column (3) of Table 16. Both series use matched within-year 2nd–98th percentile samples and raw spending. Initial taste intercepts match 1999 shares; other taste coefficients remain fixed, and calendar effects stay at their omitted reference.

tity in equation (8). In particular, T_{is} is already measured in share units and does not receive another factor of s_{is}.

To obtain B_{is}, write $s_{is} = S_{ij}q_{is}$, where S_{ij} is the household's group share and q_{is} its within-group share. At fixed utility and characteristics,

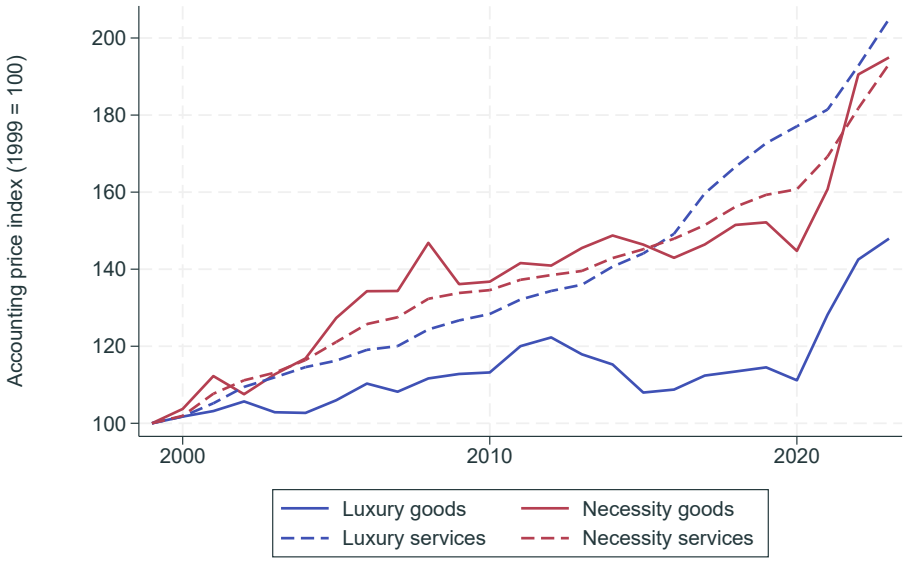
$$d\log S_{ij}\big|_{u_i,\beta_i} = (1-\sigma_0)(d\log \mathcal{P}_{ij} - d\log \mathcal{P}_i), \tag{55}$$

$$d\log q_{is}\big|_{u_i,\beta_i} = (1-\sigma_j)(d\log p_s - d\log \mathcal{P}_{ij}). \tag{56}$$

Adding the two changes gives equation (9). The budget response uses $\partial \log s_{is}/\partial \log x_i = \eta_{is} - 1$ and subtracts the expenditure change needed to compensate household-basket inflation. Composition is evaluated separately at fixed prices and budget, including the utility adjustment needed to keep that budget fixed as characteristics change.

For $\tau \in [0,1]$, each annual link linearly interpolates log prices and log budgets between its endpoints. I also interpolate the logits of the three normalized taste weights implied by endpoint characteristics; estimated demographic and regional coefficients remain fixed. At every integration point I solve household demand again and update shares, elasticities, and expenditure weights. Let $g^q_{st}(\tau)$ denote contribution $q \in \{P, I, R, D\}$ per unit of the path parameter τ on the link from $t-1$ to t, where P, I, R, and D denote relative prices, income, expenditure

Figure 33: Price inputs to the empirical demand account



Notes: Domestic consumption-VA prices, 1999–2023; 1999 = 100. Survey-weighted geometric means of household CPI exposures on the matched trimmed accounting sample, using fixed input-output coefficients.

reweighting, and composition. Cumulative contributions in percentage points are

$$\mathcal{C}_s^q = 100 \sum_{t=2000}^{2023} \int_0^1 g_{st}^q(\tau)\, d\tau. \tag{57}$$

Numerical quadrature integrates each component until the contributions converge and reconstruct the endpoint predicted-share change. No numerical discrepancy is assigned to a channel. The unexplained component is the observed-minus-predicted change, rather than an integration residual. The channel allocation is conditional on this specified path.

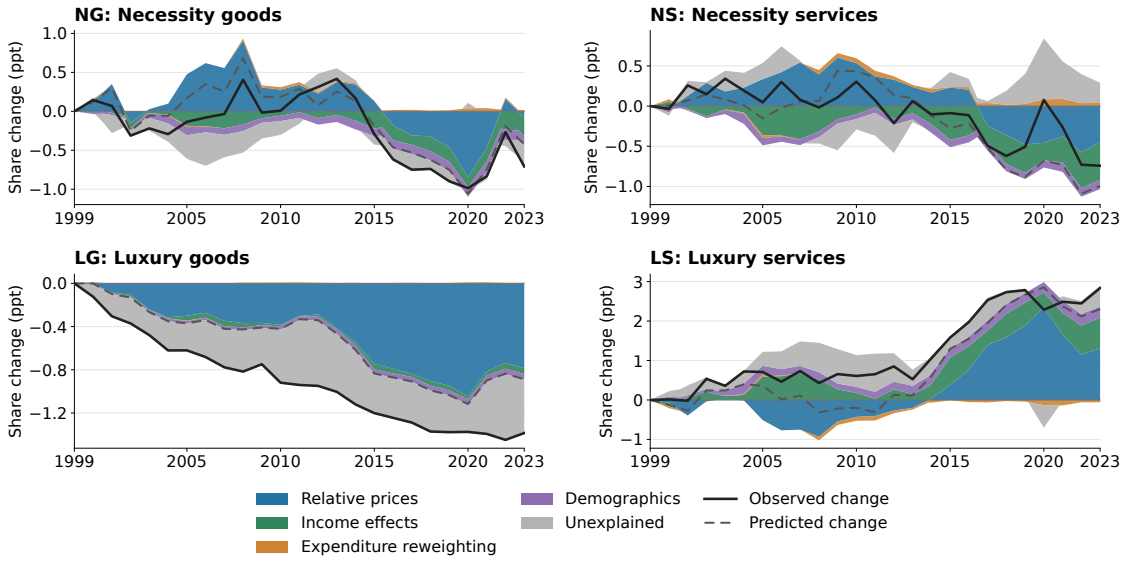
Figure 34 traces these contributions over time for all four sectors. Figure 35 subdivides the original income term along the same adjacent-year path. At each integration point, write $x_i = \mu r_i$, where $\mu = \sum_i \pi_i x_i$ is mean expenditure and $\sum_i \pi_i r_i = 1$. With population masses fixed along each rank link, $d\log\mu = \sum_i \nu_i d\log x_i$. I split the purchasing-power change in equation (54) into $d\log\mu - \sum_r s_{ir} d\log p_r$ and $d\log r_i$, applying the same household Engel responses and expenditure weights to both. Hence, the mean-expenditure component includes the entire household-basket inflation adjustment; it is not pure growth in common-CPI-deflated mean expenditure. The relative-distribution component measures changes in mean-normalized budgets through Engel responses. Expenditure reweighting remains separate: A common proportional increase in all budgets leaves expenditure weights unchanged. For luxury services, the two income components are 0.77 and 0.01 percentage points, while reweighting contributes −0.06 points.

Figure 36 instead compares finite changes in common-CPI-deflated mean expenditure and its mean-one relative distribution, holding sector prices and rank-attached household characteristics at their 1999 values. The common annual CPI is the survey-weighted geometric mean of national quarterly CPI values on the accounting population, normalized to 1999. The two-factor Shapley allocation averages the two orders of changing the mean and relative distribution. It includes

their effects on both household shares and aggregate spending weights; its endpoint need not equal the historical income contribution. This fixed-price exercise is distinct from the preceding differential accounting.

The full four-factor finite-change account additionally changes prices and household composition, evaluating sixteen combinations and averaging all 24 switching orders. Its price contribution includes purchasing-power effects relative to the common CPI, unlike the original compensated-price contribution. Rank coupling preserves population mass and each endpoint's joint expenditure and characteristic distribution. It links repeated cross-sections rather than household histories, and relative-distribution changes need not be a mean-preserving spread. Estimated demographic and regional coefficients remain fixed in all these exercises.

Figure 34: Demand-accounting contributions over time



Notes: Cumulative changes in domestic consumption-value-added shares, 1999–2023, in percentage points; column (3), with fixed taste coefficients and calendar effects held at their omitted reference. Prices hold utility fixed; income is net of household-basket inflation. Expenditure reweighting changes households' spending weights; demographics includes region. Unexplained is observed minus predicted change.

C.6 Demand accounting with changing calendar effects

The main account holds the estimated calendar effects at their omitted reference. I now restore their values at each interview date, retaining column (3)'s demand parameters, demographic and regional coefficients, matched sample, observed budgets and prices, and the same three intercepts calibrated to 1999 shares. No taste parameter is re-estimated and no later aggregate share is targeted.

Allowing these effects to change raises the predicted 1999–2023 luxury-service share increase from 2.31 to 2.54 percentage points, compared with 2.84 points in the data (Figure 37). The unexplained increase falls from 0.53 to 0.29 points. This improvement describes the contribution of estimated dated shifts within the sample. It does not identify those shifts as preference changes: They may also absorb other common changes in demand or measurement.

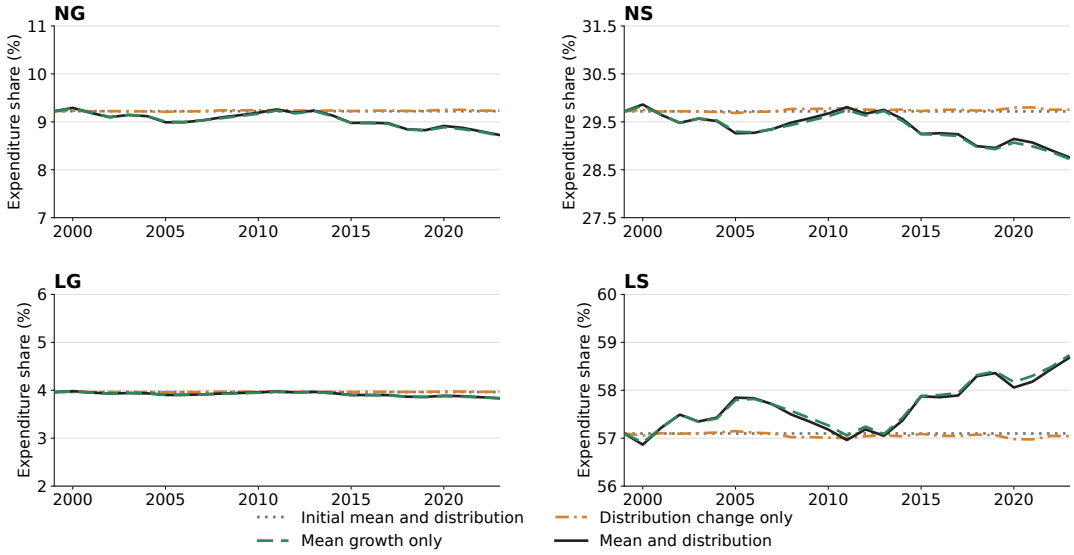
Figure 35: Average growth and distribution changes within the income contribution



Notes: Domestic consumption-value-added share changes, 1999–2023, in percentage points. The historical income contribution is split into mean expenditure, including the household-basket inflation adjustment, and relative-distribution changes through Engel responses. Expenditure reweighting is separate. Other components are unchanged from Figure 8; estimated taste coefficients remain fixed.

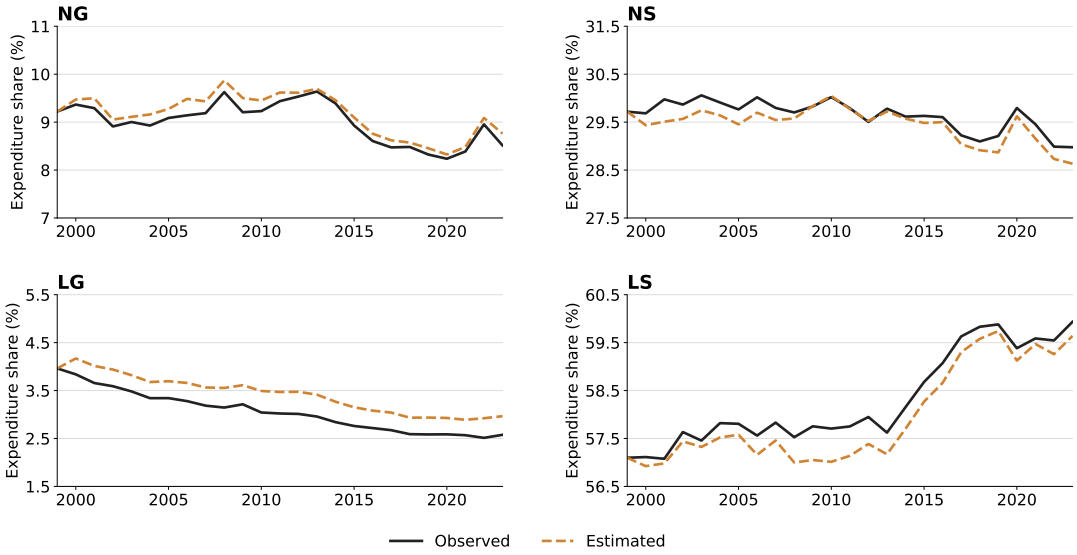
Figure 38 extends the four-term identity to include a separate calendar contribution. I integrate all terms along the same adjacent-year path with the dated effects changing. The normalization interaction between calendar effects and household composition is split equally between the two orders of changing these inputs. Calendar effects contribute 0.22 percentage points to the luxury-service expansion. Because price, income, reweighting, and composition responses are

Figure 36: Mean expenditure growth and expenditure-distribution changes



Notes: Domestic consumption-value-added shares, 1999–2023. Expenditure is deflated by a common annual CPI. Paths change its mean, its mean-one relative distribution, or both, holding sector prices and rank-attached household characteristics at 1999 values. This finite-change comparison is separate from the historical income subdivision in Figure 35.

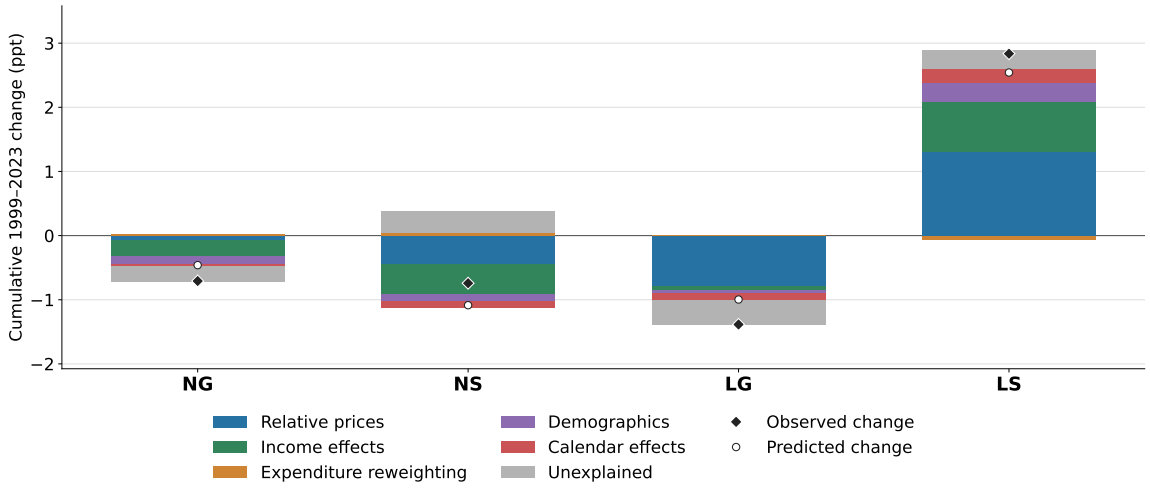
Figure 37: Observed and predicted expenditure shares with changing calendar effects



Notes: CEX, 1999–2023; domestic consumption value added. Specification (3), with estimated calendar effects restored by interview date and the same 1999 intercept calibration as Figure 7. Other estimated coefficients remain fixed. Both series use matched within-year 2nd–98th percentile samples and raw expenditure.

evaluated along the new joint path, their contributions also change slightly; the increase in the total prediction is 0.24 points. The unexplained term remains observed minus predicted change and absorbs no numerical integration discrepancy.

Figure 38: Sources of expenditure reallocation with changing calendar effects



Notes: Domestic consumption-value-added share changes, 1999–2023, in percentage points. Specification (3); estimated calendar effects vary by interview date. All contributions are recalculated jointly with the same initial-share calibration as the main account. Calendar effects and composition share their normalization interaction symmetrically. Unexplained is observed minus predicted change.

D Model details

D.1 Unit costs and factor demand

For J fixed occupation types and elasticity $\eta \neq 1$, the single labor aggregate and its unit cost are

$$\ell_{st} = \left[\sum_{j=1}^{J} a_{sjt}(\psi_t n_{sjt})^{(\eta-1)/\eta}\right]^{\eta/(\eta-1)}, \qquad \sum_j a_{sjt} = 1, \tag{58}$$

$$W_{st} = \left[\sum_j a_{sjt}^{\eta} v_{jt}^{1-\eta}\right]^{1/(1-\eta)}, \qquad v_{jt} = w_{jt}/\psi_t. \tag{59}$$

The baseline has $J = 5$ and $\eta = 1$, so $\ell_{st} = \prod_j (\psi_t n_{sjt})^{a_{sjt}}$ and $W_{st} = \prod_j (v_{jt}/a_{sjt})^{a_{sjt}}$. There is no education composite or separate education-demand shift. For general η, type j's share of the sector's labor bill is

$$\phi_{sjt} = \frac{a_{sjt}^{\eta} v_{jt}^{1-\eta}}{\sum_h a_{sht}^{\eta} v_{ht}^{1-\eta}}, \qquad w_{jt} n_{sjt} = (1-\alpha_{st})\phi_{sjt} X_{st}. \tag{60}$$

At unit substitution, $\phi_{sjt} = a_{sjt}$, yielding equation (17). Hourly wages equal service-unit returns times quality; sectoral average pay divides the wage bill by actual hours.

Initial calibration. Initial nominal value-added shares S_{s0} use the four classified KLEMS sectors. Capital shares are capital compensation divided by capital plus labor compensation. Initial ORG type pay determines relative hourly wages, with one common scale reconciling the aggregate labor bill. Initial quality is one, so hourly wages equal effective-unit returns. Let H_{sj0}^{CPS} be Basic Monthly actual hours. The seed $B_{sj0}^{\text{seed}} = w_{j0} H_{sj0}^{\text{CPS}}$ is reconciled by a positive Kullback–Leibler projection with margins

$$\sum_j B_{sj0}^{*} = (1-\alpha_{s0}) X_0 S_{s0}, \qquad \sum_s B_{sj0}^{*} = w_{j0} \bar{n}_{j0}. \tag{61}$$

Rows use KLEMS sector labor bills at model expenditure levels; columns use CEX-derived type supplies and rescaled ORG wages. At unit substitution, $a_{sj0} = B_{sj0}^{*} / \sum_h B_{sh0}^{*}$. Initial tastes match S_{s0} and Hicks-neutral levels normalize prices to one. These choices reproduce every cell of the reconciled bill matrix. Initial model hours need not equal the raw CPS cells, because balancing changes the seed. Later sectoral allocations remain endogenous conditional on aggregate supplies and calibrated relative-wage changes.

Common quality. CPS type-hour growth is rebased on the initial CEX supplies and normalized to aggregate hours per household. The reference compensation weights use those hours and the composition-adjusted ORG wage paths. I integrate equation (24) along log-linear annual hours and compensation paths. Since weights sum to one, adding the same $d \log \psi_t$ to every type's hour growth exactly reconciles their compensation-weighted services index with KLEMS. This

avoids adding an aggregate labor-quality adjustment and counting the occupation-composition contribution a second time.

D.2 Household assignments and expenditure financing

The ownership distribution starts from 2,496 CEX households in 1997. I retain their observed member hours, expenditure, survey weights, and the previously constructed capital match from the 1998 SCF, whose income measure refers to 1997. SCF donors have a reference person aged at least 25. Matching takes place within pretax-income decile and four education categories: Less than high school, high-school graduate, some college or an associate degree, and a bachelor's degree or higher. Within each cell, I select the 20 closest distinct donor families (or all available families if fewer) using age, economic-unit size, income rank, and homeownership. The distance averages squared differences in age divided by ten, unit size, and income rank divided by 0.1, together with an indicator for different homeownership status. A donor and imputation are drawn using normalized SCF survey weights. CEX asset reports do not enter this match.

All five SCF imputations enter with equal weight. Investment holdings are floored at zero and normalized to mean one within each imputation before combining them. The draw provides a conditional ownership rank. I then align these ranks to the eligible SCF distribution of full investment holdings while preserving the original CEX weights. In particular, order CEX households by their matched ownership scores, let F_i^- and F_i^+ be the endpoints of household i's interval in the weighted distribution, and let $Q^{\mathrm{SCF}}(u)$ denote the quantile function of the mixture of mean-one SCF holdings. The aligned score is

$$k_i = \frac{1}{\pi_i} \int_{F_i^-}^{F_i^+} Q^{\mathrm{SCF}}(u)\, du, \qquad F_i^+ - F_i^- = \pi_i. \tag{62}$$

This preserves the SCF marginal ownership distribution at the resolution of CEX household weights. It does not observe consumption and wealth jointly or identify their true conditional dependence. Income ranks determine the matching cells; the expenditure ranks used to report ownership and welfare are calculated separately from initial CEX consumption-value-added expenditure divided by the square root of household size.

Capital endowments are $\bar{K}_{it} = \bar{K}_t k_i / (\sum_h \pi_h k_h)$. Thus, SCF determines relative ownership and KLEMS determines aggregate capital services. Education remains in this fixed wealth match; it is absent from occupation-type definitions, labor matching, production, and baseline labor-service measurement. I keep the capital match fixed across the occupation and elasticity comparisons.

Worker assignments use 1998 CPS ASEC donors, whose hours and earnings refer to 1997. The pooled matching uses observed occupation, hours, employment status, employee pay, age, and sex, without education. Joint assignments preserve dependence between household members. Every replica retains its parent's capital, total hours, and initial expenditure; replica weights sum to the parent's survey mass and reproduce aggregate type supplies. The resulting 13,830

replicas and their weights remain fixed across experiments. These are imputed joint portfolios, not directly observed detailed occupation holdings in CEX.

Let x_{i0} be initial expenditure and y^F_{i0} factor income. A common scale $\kappa = (\sum_i \pi_i y^F_{i0})/(\sum_i \pi_i x_{i0})$ gives $y_{i0} = \kappa x_{i0}$ and balanced residuals $T_{i0} = y_{i0} - y^F_{i0}$. Payer fractions are $\tau_i = \max(-T_{i0}, 0)/y^F_{i0}$ when factor income is positive and zero otherwise. Recipient coefficients are $b_i = \max(T_{i0}, 0)/\mathcal{R}_0$, with $\mathcal{R}_0 = \sum_i \pi_i \max(T_{i0}, 0)$. The fixed coefficients preserve positive budgets and balanced financing, including for households with no measured factor holdings.

Gross initial payments equal 33.38% of scaled expenditure. This is a substantial reconciliation between measured factor income and spending, rather than an estimated tax-and-transfer system. The residual can include transfers, saving and dissaving, omitted resources, and measurement differences. Expenditure above factor income is financed through this rule rather than assigned to capital ownership.

D.3 Reference productivity with changing production weights

Let $\mu_s(\mathbf{v}, \mathbf{z})$ denote unit cost before dividing by Hicks-neutral productivity, where $\mathbf{z}$ contains the capital share and normalized occupation weights. Their contribution to productivity growth, holding effective factor returns fixed, is $-\nabla_{\mathbf{z}} \log \mu_s \, d\mathbf{z}$. I recover Hicks-neutral growth from KLEMS reference TFP as

$$\log \frac{A_{st}}{A_{s0}} = \log \frac{T^{\text{ref}}_{st}}{T^{\text{ref}}_{s0}} - \Gamma^{\text{ref}}_{st}, \qquad \Gamma^{\text{ref}}_{st} = -\int_0^t \nabla_{\mathbf{z}} \log \mu_s(\mathbf{v}^{\text{ref}}_u, \mathbf{z}_{su}) \, d\mathbf{z}_{su}. \tag{63}$$

The reference return path is obtained in the historical calibration and held fixed across counterfactuals, together with the production-weight paths and their corrections. Realized productivity evaluates the same changes at equilibrium returns. Equal reference TFP growth therefore need not equalize Hicks-neutral or realized productivity growth.

Reporting prices. I chain a Divisia price index along each economy's equilibrium path. Between annual dates, log reference TFP, log supplies and quality, capital shares, and unnormalized log type weights are interpolated linearly. Weights are normalized at each step and Hicks-neutral productivity includes the corresponding partial reference correction. Trapezoidal integration uses endogenous expenditure shares and price changes, with numerical refinement until equilibrium and accounting tolerances are met. These weights differ from the observed annual midpoint weights defining the common-growth experiment.

Let $k = 0, \ldots, M_t$ index the equilibrium steps within annual interval t, and let $S_{s,t,k}$ be the corresponding sectoral expenditure shares. The reporting index is

$$\log \frac{P_t}{P_{t-1}} = \sum_{k=1}^{M_t} \sum_s \frac{S_{s,t,k-1} + S_{s,t,k}}{2} \log \frac{p_{s,t,k}}{p_{s,t,k-1}}, \qquad P_0 = 100. \tag{64}$$

For the income–price attribution, I multiply each economy's prices and budgets by the common CPI-U index divided by its own Divisia index, both initially 100. A common rescaling leaves

exact welfare unchanged, but this convention fixes the units of hybrid valuations that combine one economy's prices and another's budget. The channel decomposition is conditional on this normalization.

D.4 Matched homothetic preferences and welfare

The matched expenditure function in equation (26) reproduces initial expenditure, Hicksian demands, and compensated price derivatives household by household. At fixed prices, its demands scale proportionately with its utility index. Each counterfactual equilibrium is then valued under the initial nonhomothetic preferences using equation (28). Reoptimization for this valuation does not change the equilibrium.

Deciles use original CEX population weights and initial consumption-VA expenditure divided by the square root of household size. Boundary mass is split, with ties shared proportionately. Replicas inherit their parent's rank; groups remain fixed across scenarios. Size adjustment changes the ranking, rather than weights or the initial total-budget denominator. Dispersion is calculated across joint replicas, without first averaging each parent's assignments.

Define household i's reference money metric as

$$M_i(\mathbf{p}, y) = \mathcal{E}_i^0\Big(\mathbf{p}_0, V_i^0(\mathbf{p}, y)\Big).$$

Suppressing time subscripts, the income contribution between counterfactual a and baseline b is

$$\frac{1}{2}\{M_i(\mathbf{p}^a, y_i^b) - M_i(\mathbf{p}^a, y_i^a) + M_i(\mathbf{p}^b, y_i^b) - M_i(\mathbf{p}^b, y_i^a)\}. \tag{65}$$

The price contribution switches prices at each income endpoint and averages the two orders. Their sum is the exact money-metric welfare difference. Reported contributions are expressed as percentages of the initial budget by multiplying these money-metric differences by $100/y_{i0}$. For the nonhomothetic comparison, I further split income into capital and all labor by averaging their two switching orders at both price endpoints. Effective factor claims include the associated financing flows. There is no college–non-college attribution in this model.

Factor exposure and real income. Under baseline financing, household budgets are linear in factor returns and aggregate factor quantities. Let $\widetilde{F}_{ift}$ denote the household's factor claims after financing and $e_{ift} = v_{ft}\widetilde{F}_{ift}/y_{it}$ their income shares. With fixed relative ownership and financing coefficients,

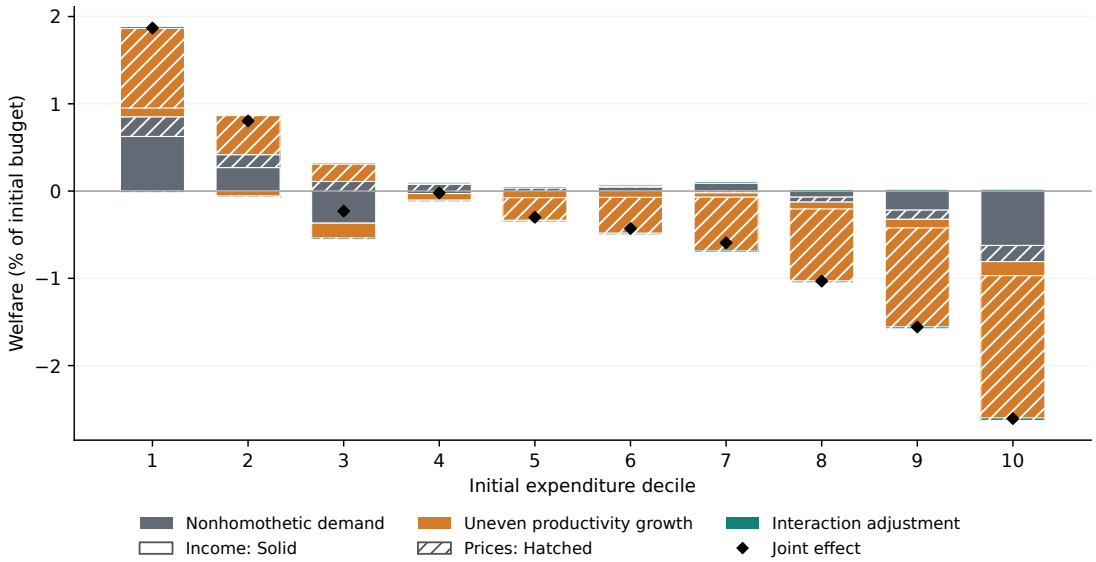
$$d\log y_{it} - \sum_s s_{ist}\, d\log p_{st} = \sum_f e_{ift}\, d\log v_{ft} + \sum_f e_{ift}\, d\log \bar{F}_{ft} - \sum_s s_{ist}\, d\log p_{st}. \tag{66}$$

The first term transmits changes in wages and capital returns according to household exposure. The second records changes in the quantities of capital and labor services, and the final term deducts inflation in the household's own basket. Integrating this identity gives a real-income account along the equilibrium path. It differs from the finite common-preference equivalent vari-

ation used in the welfare figures; in particular, an integrated log-income gain is not a percentage of the initial budget.

The interaction between nonhomothetic demand and uneven productivity growth can also be allocated across income and prices. For each channel, subtract its two separate-comparison contributions from its joint-comparison contribution. The two resulting channel interactions add to the interaction in Figure 12; no additional welfare component is left unallocated. Figure 39 displays the resulting six contributions in a single account.

Figure 39: Nested income and price contributions to welfare



Notes: 2024 welfare differences, in percent of initial total budgets, averaged within fixed initial size-adjusted expenditure deciles. Colors distinguish nonhomothetic demand, uneven productivity growth, and their interaction. Solid portions are income contributions; hatched portions are price contributions. Within each channel, the interaction equals the joint contribution minus the two separate contributions. All six portions sum to the joint welfare effect, shown by the black diamond. Common initial preferences, prices, and budget denominators apply throughout.

E Additional quantitative results

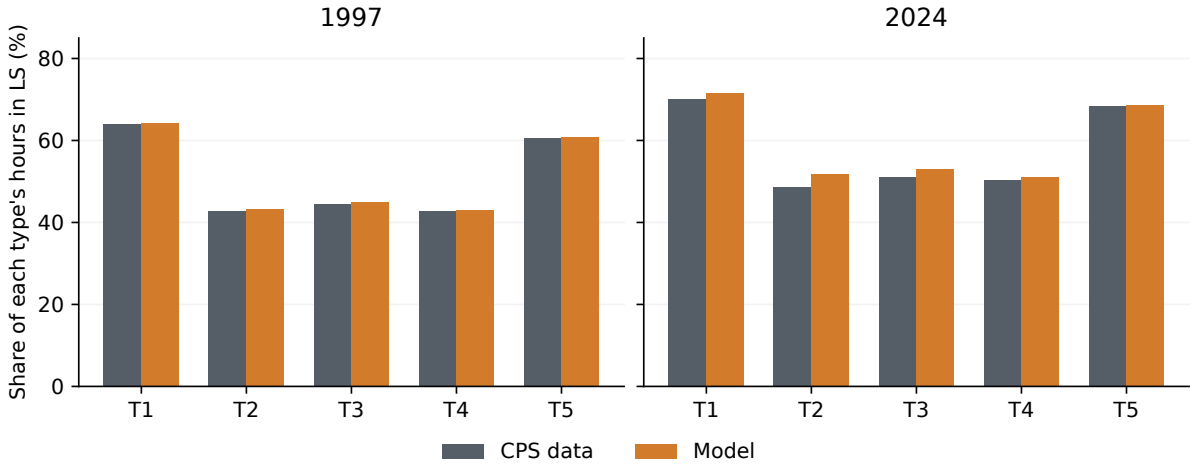
E.1 Additional model validation

The additional fit comparisons distinguish calibration targets from outcomes. Figure 40 compares the share of each occupation group's hours worked in luxury services in the CPS and the model. The initial comparison shows how closely labor-bill reconciliation preserves observed exposure; the 2024 comparison assesses the later allocation.

Figure 41 gives the more demanding comparison of changes in each occupation's share of hours within each sector. These changes are not targeted by the common annual demand shifts.

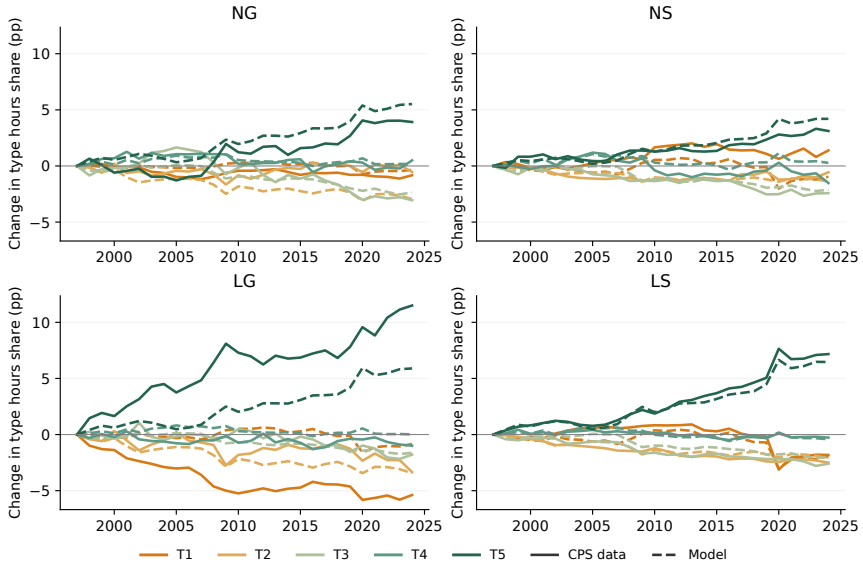
Matching relative wages does not imply matching absolute real wage growth. Figure 42 shows aggregate arithmetic mean pay and the five composition-adjusted type wage indexes. The model overstates real wage growth for the fixed types, despite matching their relative movements. The

Figure 40: Luxury-service exposure by occupation type: Data and model



Notes: CPS Basic Monthly and model, 1997 and 2024. LS hours divided by classified NG, NS, LG, and LS hours within the same occupation type and year. CPS private employed workers aged 16 and above; survey-weighted actual main-job hours, averaged over all months. Types are fixed pooled occupation groups ranked by initial pay. Initial hours enter calibration; later allocation is untargeted.

Figure 41: Changes in sector-by-occupation hours shares: Data and model



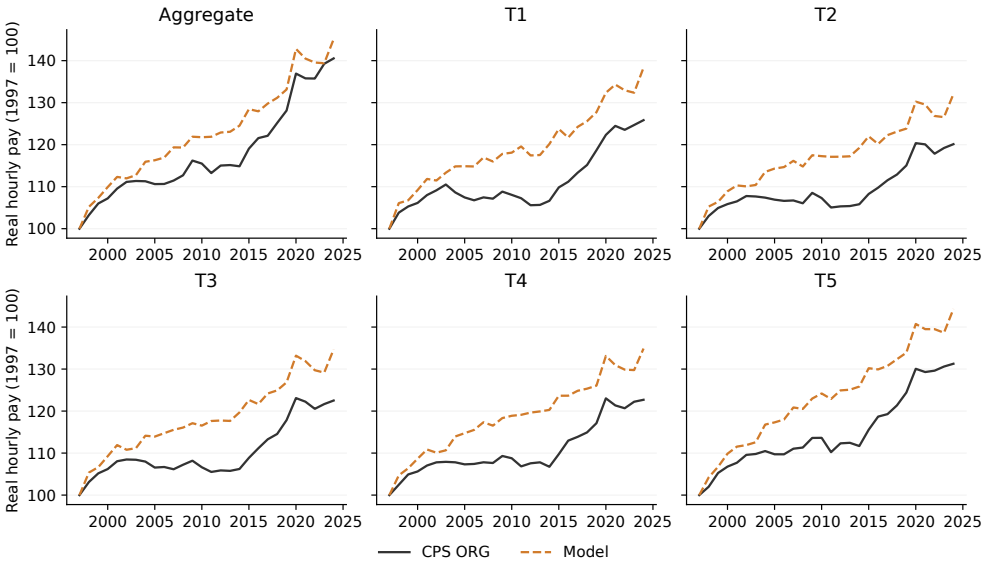
Notes: Changes from 1997 in each type's share of the sector's classified hours, in percentage points. CPS Basic Monthly actual hours and model actual hours, 1997–2024. Colors identify the five fixed occupation groups. Solid: Data. Dashed: Model. All panels have the same vertical scale.

aggregate comparison also incorporates changes in occupation composition, so it need not have the same discrepancy.

Figure 43 combines the wage and allocation implications at the sector level. Data use CPI-U and usual hours among ORG employees, whereas model sector averages use actual hours and the aggregate Divisia deflator. These sample and deflator differences matter for interpreting the fit.

Figure 44 separately reports the four targeted relative wage paths.

Figure 42: Aggregate and occupation real hourly pay: Data and model

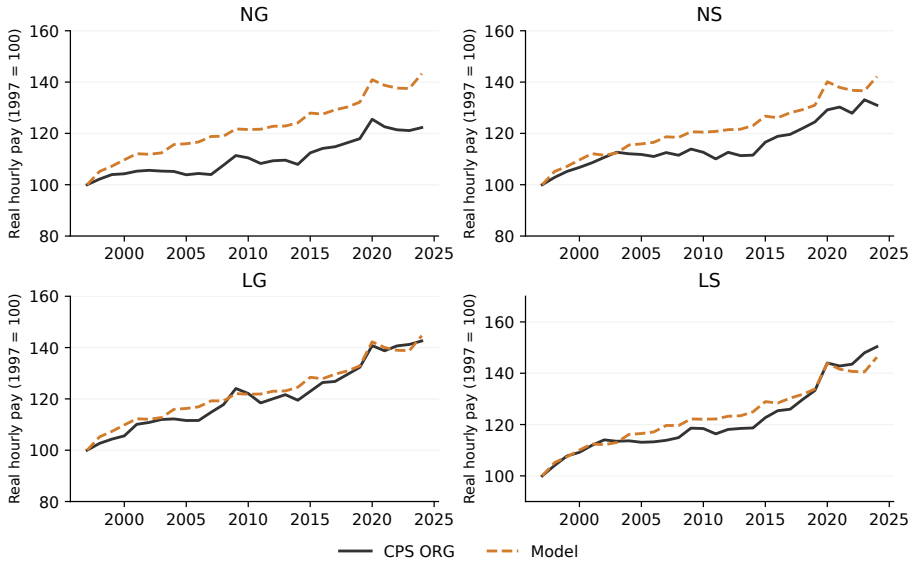


Notes: Indexes, 1997 equals 100. Aggregate pay is the current-composition arithmetic mean, calculated as the wage bill divided by hours. T1–T5 data are composition-adjusted ORG wage indexes for the fixed pooled groups, including the corrected 2003 link. Data use CPI-U; model hourly wages include common labor quality and use the model Divisia deflator. Absolute real wage growth is not a calibration target.

E.2 Demand accounting along the model baseline

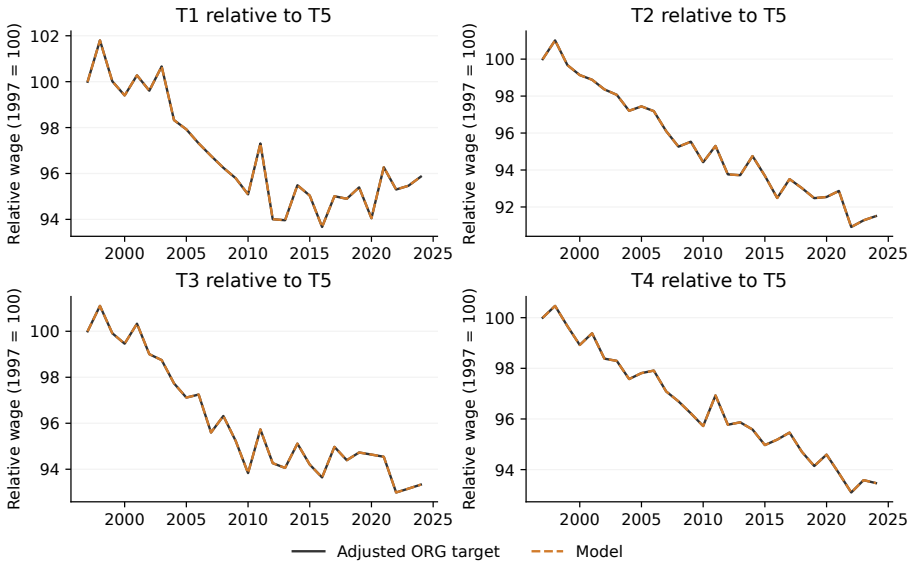
Figure 45 applies the demand account to the model's historical baseline. I interpolate log household budgets and sectoral prices between annual equilibria, keeping household identities, population weights, and tastes fixed. Compensated prices describe substitution at fixed utility.

Figure 43: Sectoral real hourly pay: Data and model



Notes: Arithmetic sectoral wage bills divided by hours; each series rebased to 100 in 1997. CPS ORG private employees aged 16–64, earnings weights times usual hours and classified sector exposure; CPI-U deflation. Model averages use actual hours and the own-economy Divisia index. Workforce composition changes over time. All sector panels have equal vertical spans and tick spacing.

Figure 44: Targeted relative wages of occupation types



Notes: Composition-adjusted ORG and model hourly wages relative to T5, rebased to 100 in 1997. Four common annual type-demand shifts target these four relative growth paths. This is a calibration check, distinct from the absolute real wage comparisons.

The Engel term is split into mean expenditure growth net of household-basket inflation and changes in relative household expenditure. A fourth term records changes in households' shares of total spending. All four terms are integrated along the same path and sum to the change in aggregate expenditure shares.

This connects the estimated demand system to the model's implied budgets and prices. It is an accounting exercise along one path, rather than an attribution to exogenous productivity shocks or a welfare decomposition. It covers producer value added over 1997–2024; the empirical demand account uses domestic consumption value added over 1999–2023. Hence, their numerical contributions need not coincide.

E.3 Counterfactual allocations and factor returns

The following figures compare the same four economies as the main welfare exercise.

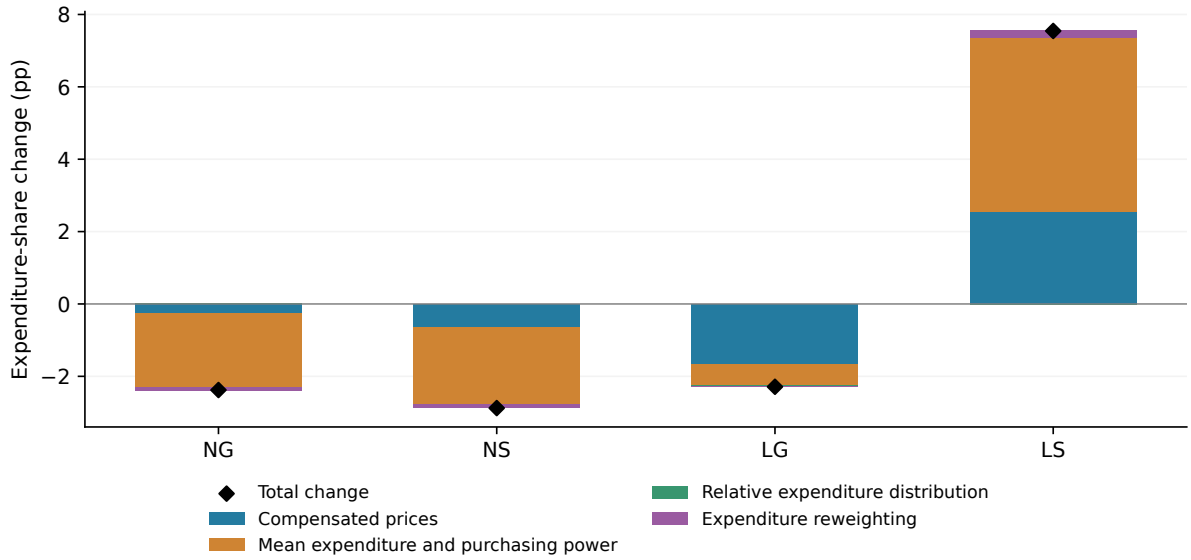
Type wages refer to fixed occupation groups. Sectoral average pay also reflects endogenous changes in the mix of types employed, so it is a different object.

Figures 49 and 50 show percentage differences from the historical baseline to make the counterfactual gaps visible. Table 17 reports the 2024 real returns to all five labor types and capital across the three comparisons.

E.4 Historical welfare gains and household heterogeneity

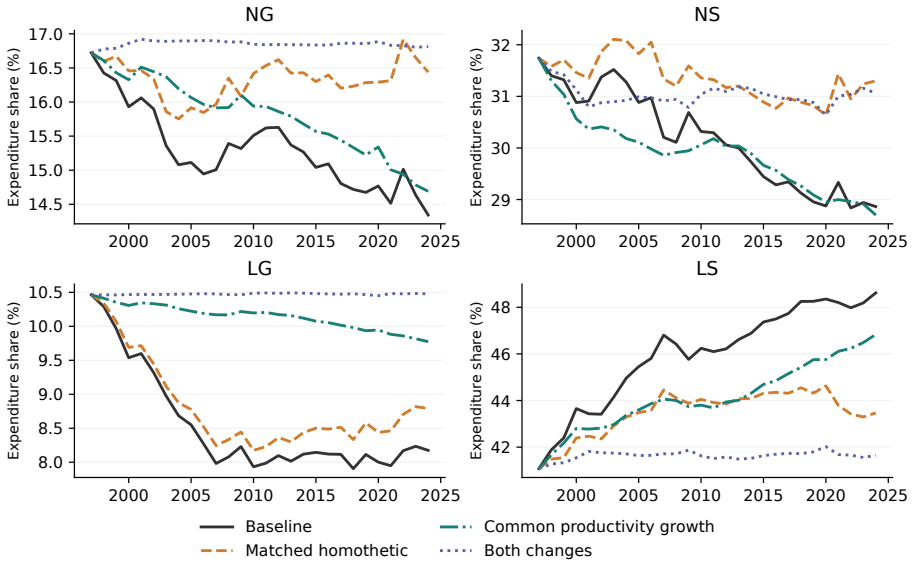
The main figures compare the historical baseline with counterfactual economies in 2024. A loss in that comparison does not imply a loss relative to 1997. Figure 51 therefore reports overall

Figure 45: Demand accounting along the model baseline



Notes: Changes in sectoral expenditure shares, 1997–2024, in percentage points. All components use the same interpolated annual equilibrium budgets and prices, with fixed household population and tastes. Diamonds show the total change. Numerical integration verifies adding up; no unexplained residual is assigned to a component.

Figure 46: Sectoral expenditure shares across counterfactuals

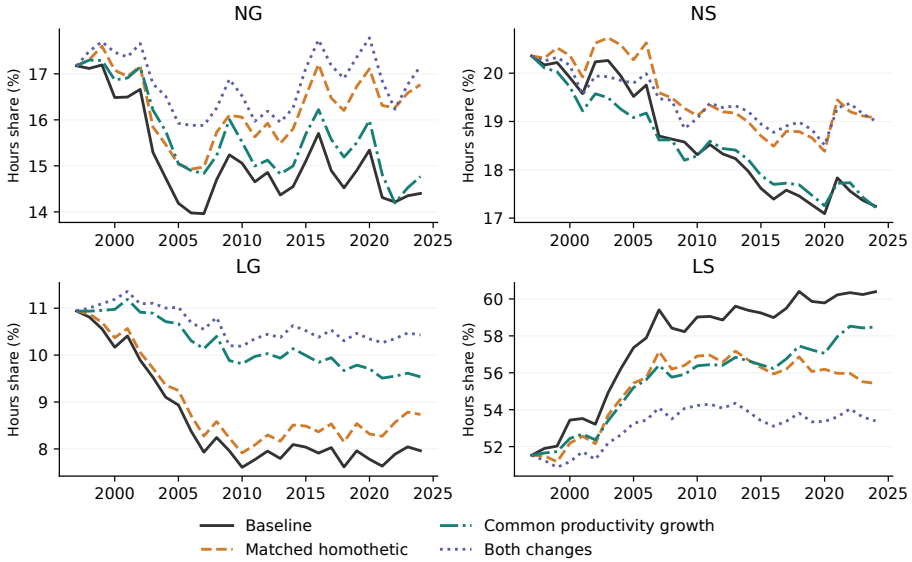


Notes: Four-sector producer-VA expenditure shares, 1997–2024. All scenarios retain baseline factor supplies, quality, capital shares, and occupation-weight paths.

gains since the initial year under the same reference preferences and prices. Factor quantities are exogenous, and these gains do not deduct a utility cost of supplying more hours.

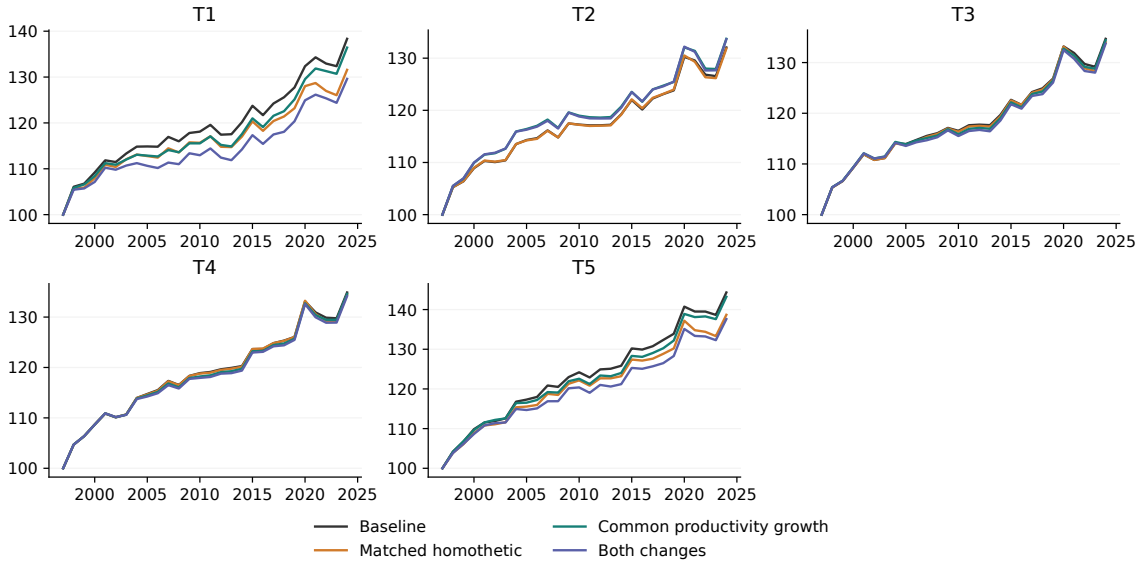
Figure 52 shows the evolution of the counterfactual gaps and their average over the 27 years from 1998 through 2024. The average gives every year equal weight and uses the same initial budget denominator throughout. It is neither a growth rate nor lifetime equivalent variation.

Figure 47: Sectoral hours shares across counterfactuals



Notes: Actual hours as a share of the four-sector total. Worker types remain fixed and can move across sectors.

Figure 48: Real hourly wages of fixed occupation types

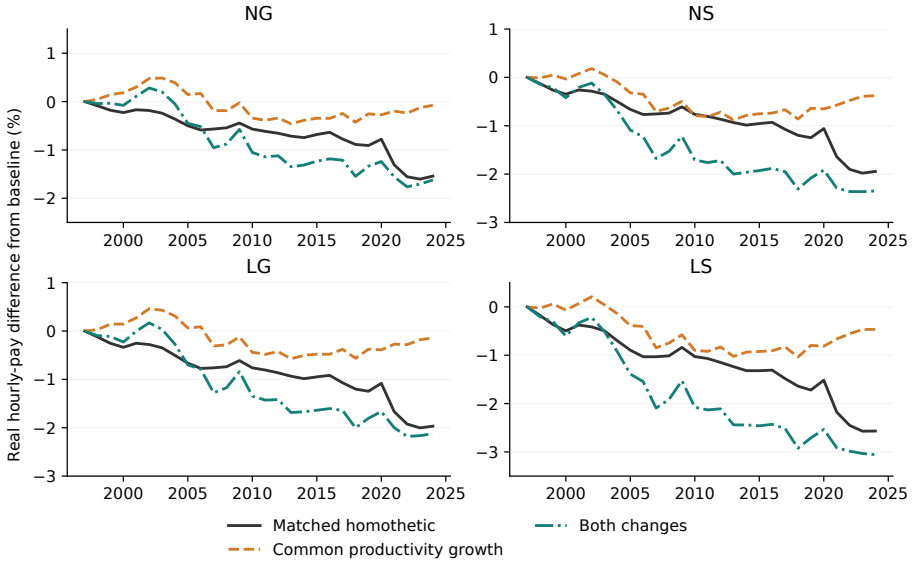


Notes: T1–T5 are fixed occupation groups ranked by initial occupation pay. Hourly wages include common labor quality; each economy uses its own Divisia deflator. Indexes equal 100 in 1997.

Figure 53 shows who lies behind within-decile dispersion. Each point represents one of the 2,496 original CEX households, averaging welfare over its fixed occupation-assignment replicas. The figure distinguishes households supplying labor from those with no measured labor supply. This averaging is useful for relating outcomes to observed households, but removes assignment uncertainty within a parent household. The percentile ranges in Figure 16 and the group comparisons below retain the joint replicas.

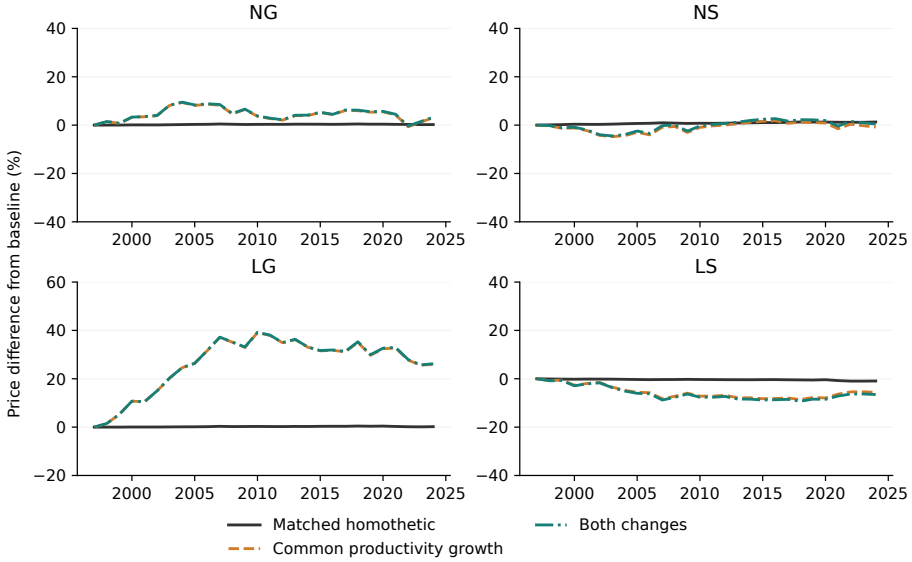
Figure 54 compares the two separate welfare effects for the same households. Their population-weighted correlation is 0.57. Households gaining from both account for 25.15% of the population,

Figure 49: Counterfactual sectoral real pay relative to baseline



Notes: 100 × (counterfactual/baseline − 1). Sectoral labor bills divided by actual hours, deflated by each economy's aggregate Divisia index. Changes reflect fixed-type wages and workforce composition. Equal vertical spans and tick spacing across sectors.

Figure 50: Counterfactual sectoral prices relative to baseline



Notes: 100 × (counterfactual/baseline − 1). Prices are relative to each economy's aggregate Divisia index. Equal vertical spans and tick spacing across sectors.

while 44.02% lose from both. The remaining 30.83% gain in one comparison and lose in the other. Thus, the two mechanisms often favor the same households, but the correspondence is far from complete.

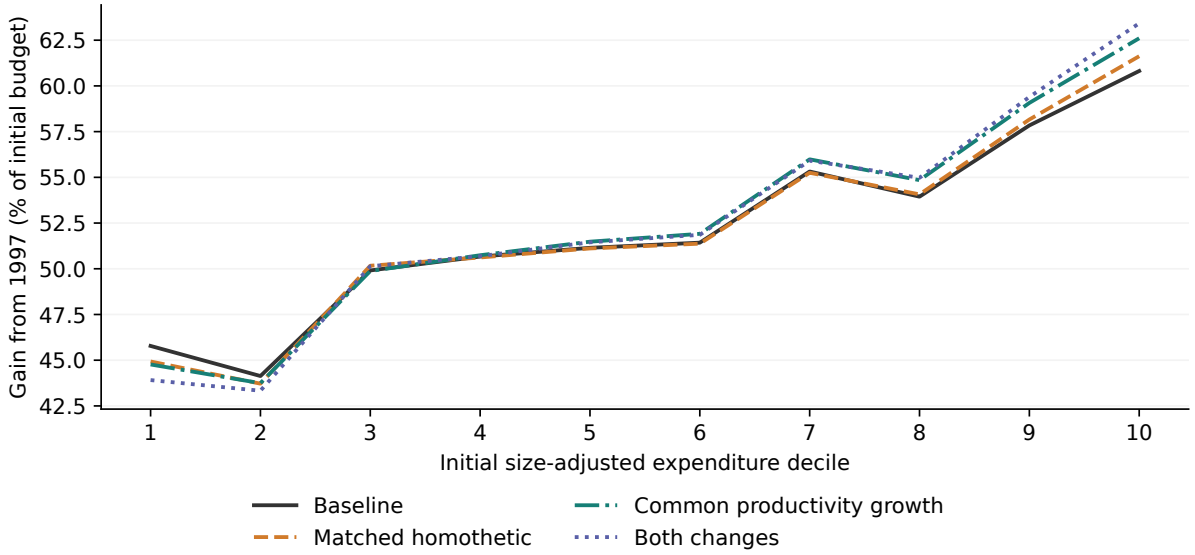
Figure 55 reports labor supply, reference-person age, sex, marital status and education, household size, and housing tenure. Filled diamonds show group means. Open circles give each original expenditure decile a weight of 10% within each group, separating differences in expenditure composition from differences among households at comparable ranks. This adjustment

Table 17: Real factor-return gains across counterfactual comparisons, 2024

Factor	Nonhomothetic	Uneven growth	Both
T1	5.14	1.40	6.74
T2	0.16	-1.24	-1.18
T3	0.51	0.21	0.72
T4	0.31	0.21	0.54
T5	4.06	0.76	4.88
Capital	-2.68	-0.57	-3.18

Notes: 100 × (baseline/counterfactual − 1), using each economy's aggregate Divisia deflator. Columns remove nonhomotheticity, unequal reference-productivity growth, or both. Positive values mean a higher real return in the historical baseline. Capital is the rental rate per unit of services.

Figure 51: Overall welfare gains since 1997



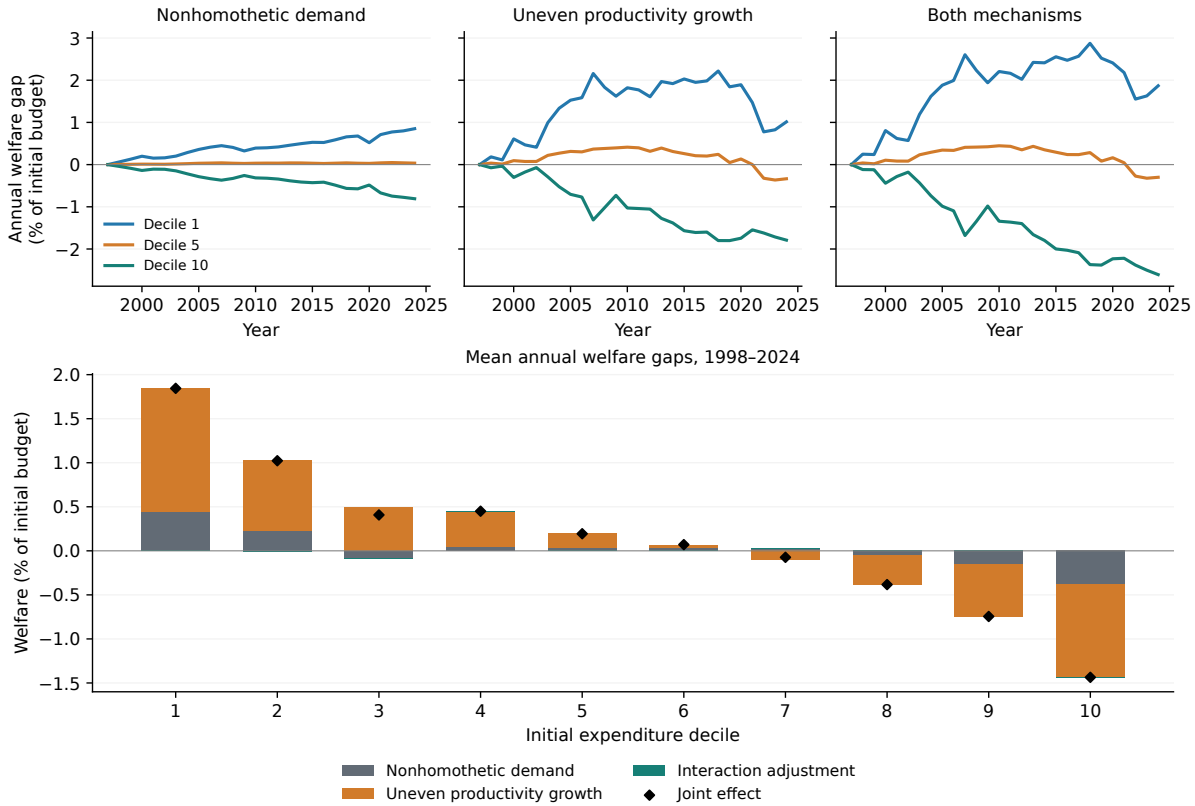
Notes: 2024 budget opportunities valued under initial nonhomothetic preferences and prices, relative to the same household's 1997 budget. Means within fixed initial size-adjusted expenditure deciles. These are historical gains in each economy, rather than baseline-minus-counterfactual effects.

is unavailable when a group has no observations in a decile. Intervals show the 10th–90th percentiles of model-household outcomes, not confidence intervals. These are descriptive comparisons of initial household groups, not causal effects of demographic characteristics or life-cycle predictions. Homeownership also enters the SCF wealth match. Reporting education groups here does not change occupation definitions or introduce education into labor demand.

E.5 Sectoral productivity experiments and Shapley accounting

Figure 56 gives the hours counterpart to Figure 17. Equalizing one sector's reference-productivity growth to the annual four-sector weighted mean can change both the composition and the weighted mean of progress. Figure 57 holds the latter fixed. After setting the named sector's growth to the same benchmark, I add a common adjustment to the other three sectors so that the observed-share-weighted mean still equals baseline growth in every year. Table 18 reports LS hours expansion and its reduction relative to baseline under both designs. Fixing this reference-

Figure 52: Welfare gaps over time and their annual averages



Notes: Top: Annual baseline-minus-counterfactual welfare for initial expenditure deciles 1, 5, and 10. Bottom: Equally weighted mean annual gaps over 1998–2024 for all ten deciles. Bars show nonhomothetic demand, uneven productivity growth, and the interaction adjustment; diamonds show their sum, the joint effect. The interaction is joint minus the two separate effects, averaged over the same years. Every observation uses common initial nonhomothetic preferences, initial prices, and the household's initial annual budget. Units: Percent of that budget.

growth average does not fix household purchasing power, Hicks-neutral productivity, or realized aggregate productivity.

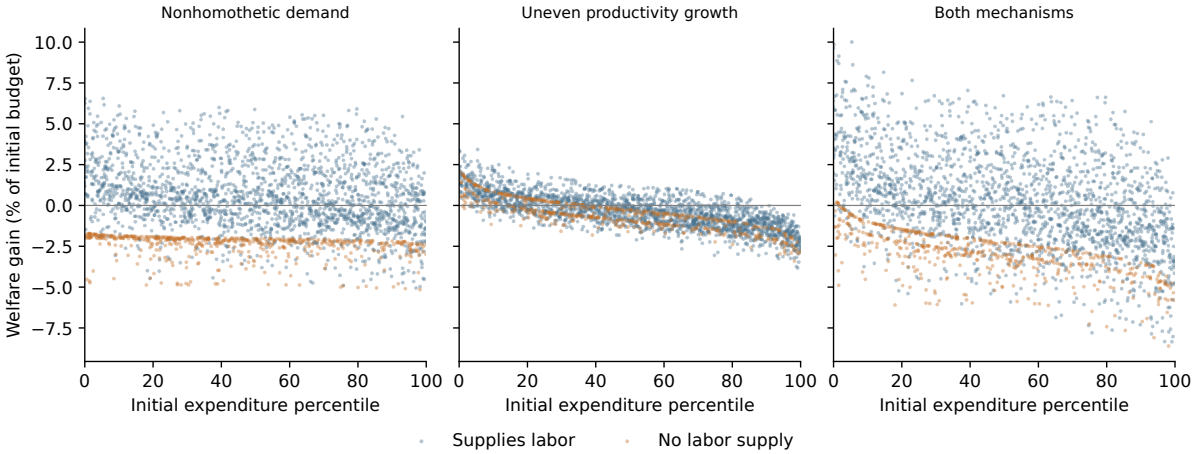
Table 18: LS hours expansion with and without preserving average reference growth

	No offset elsewhere		Mean growth preserved	
Growth equalized	LS change	Reduction	LS change	Reduction
Baseline	8.86	0.00	8.86	0.00
NG	8.67	0.20	8.69	0.17
NS	8.96	-0.10	8.95	-0.09
LG	7.64	1.23	7.77	1.10
LS	8.29	0.57	7.42	1.44

Notes: Percentage-point changes from 1997 to 2024. First pair: Named sector equalized to the four-sector benchmark, with no adjustment elsewhere. Second pair: The other sectors adjust to preserve annual weighted-mean reference growth. Reduction equals baseline expansion minus the experiment's expansion; a negative reduction means more expansion.

Both luxury sectors' distinctive growth paths support LS labor absorption in these experiments, but their ranking depends on the benchmark. Without an offset elsewhere, equalizing LG growth reduces LS hours expansion by 1.23 percentage points, compared with 0.57 for LS. Preserving

Figure 53: Household welfare across initial expenditure ranks



Notes: 2024 baseline-minus-counterfactual welfare, in percent of initial total budgets. One point per original CEX household, averaging its fixed replicas; all 2,496 households are retained. Horizontal axis: Household-weighted 1997 expenditure percentile after division by the square root of household size. Points have equal size.

mean reference growth changes these reductions to 1.10 and 1.44 points, respectively. Thus, the stronger LG effect in the first comparison does not extend to the mean-preserving experiment.

Figure 58 reverses the experiment by retaining the distinctive historical growth of LG, LS, or both, while assigning the remaining sectors their own weighted-average growth. This also preserves mean reference growth in every year. Retaining only LG generates an 8.04-percentage-point rise in the LS hours share, compared with 8.41 points when only LS retains its historical growth. Retaining both luxury sectors gives 8.80 points, close to the baseline's 8.86. This does not attribute the whole expansion to their distinctive growth: LS hours still expand by 6.95 points under common growth, with nonhomothetic preferences and the other historical inputs retained.

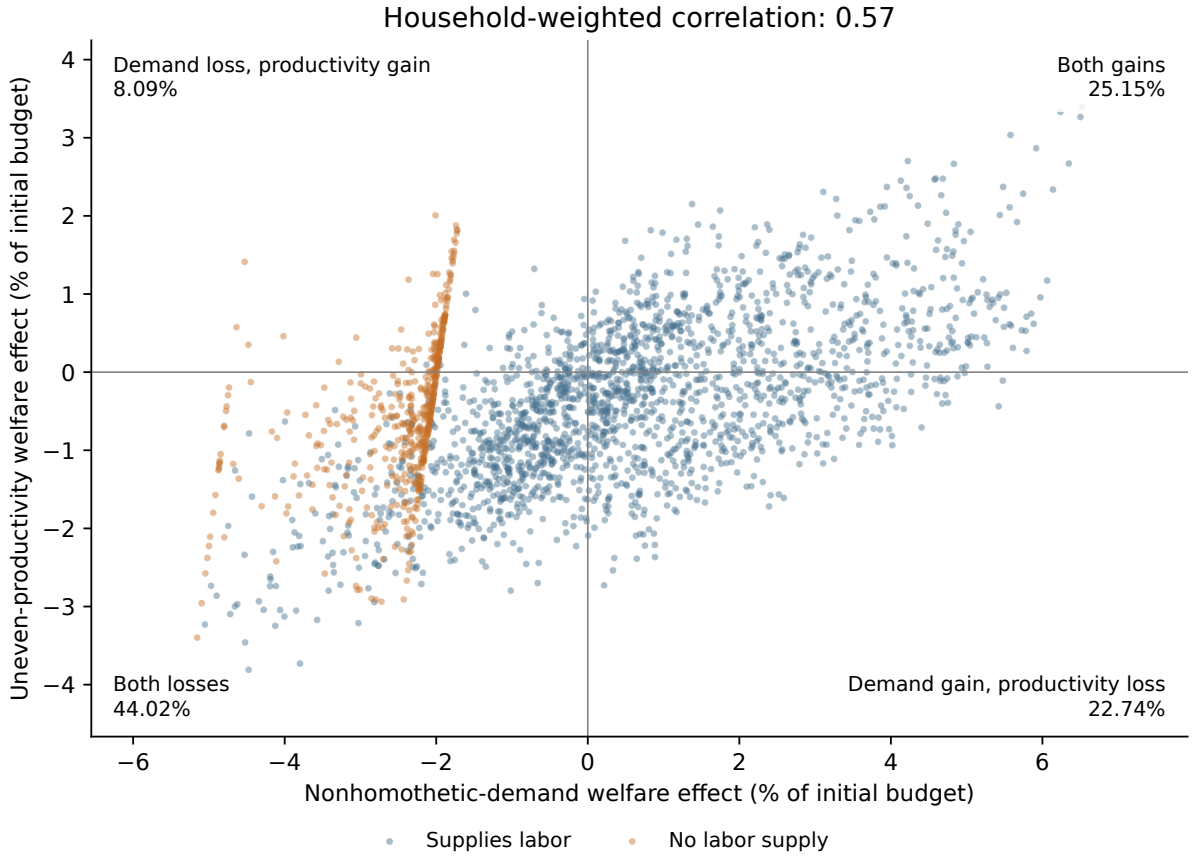
The Shapley account uses five switches: Nonhomothetic demand and each sector's distinctive reference-productivity growth. A removed productivity path is replaced by the same observed-share-weighted common-growth path in equation (27), without compensating other sectors. For each of the 32 configurations I solve the endpoint equilibrium with initial households and all other calibrated paths unchanged. The four main endpoints coincide with their historical-path solutions.

If $f(A)$ is the LS expenditure-share change with switches A active and N contains all five switches, component k receives

$$\Phi_k = \sum_{A \subseteq N \setminus \{k\}} \frac{|A|!(4-|A|)!}{5!} \big[f(A \cup \{k\}) - f(A) \big]. \tag{67}$$

The contributions plus $f(\varnothing)$ sum to baseline expansion. Averaging all 120 switching orders gives 5.17 percentage points for nonhomothetic preferences, 1.00 for LG, 0.77 for LS, 0.14 for NG, and −0.11 for NS. Figure 59 shows these means; Table 19 reports the minimum and maximum

Figure 54: Household welfare from nonhomothetic demand and uneven productivity growth



Notes: Each point is one of the 2,496 original CEX households, averaging its fixed occupation-assignment replicas. Both axes report 2024 baseline-minus-counterfactual welfare under common initial preferences and prices, in percent of the same initial total budget. Colors distinguish measured labor supply. Points have equal size; the correlation and four gain/loss quadrant shares use original household population weights. These statistics describe household-average outcomes, rather than dispersion across individual replicas.

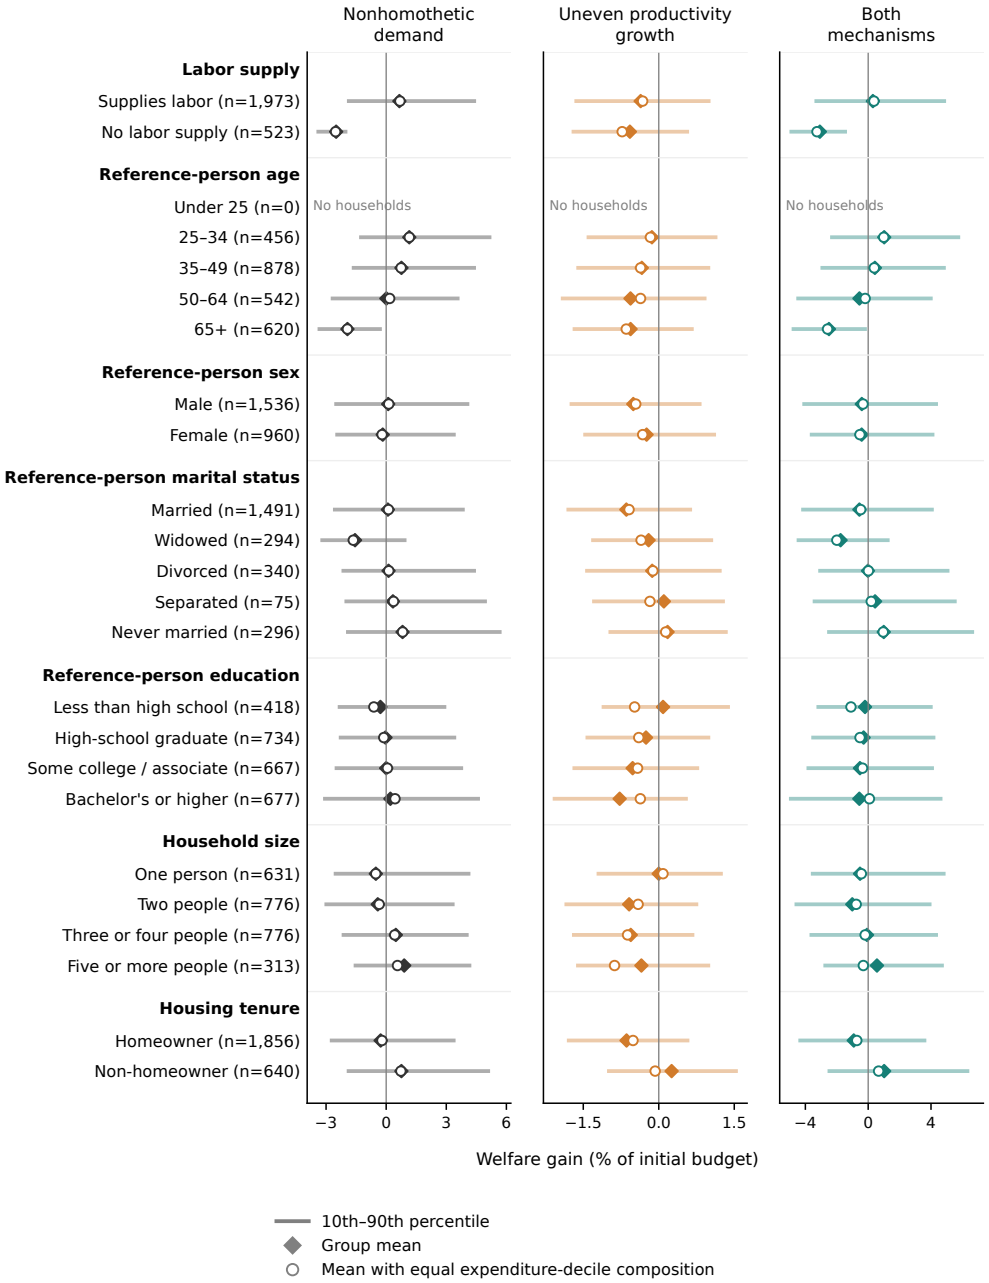
contributions across orders for every component. These ranges describe interactions and order dependence, not statistical uncertainty.

The remaining 0.57 percentage points under homothetic preferences and common growth consist of 0.45 points from household price responses and 0.12 from changes in expenditure weights. Let $\nu_{it} = \pi_i y_{it} / \sum_h \pi_h y_{ht}$ and let bars denote arithmetic averages across adjacent years. The exact product rule is

$$S_{\mathrm{LS},T} - S_{\mathrm{LS},0} = \sum_{t=1}^{T} \sum_i \bar{\nu}_{it} \Delta s_{i\mathrm{LS},t} + \sum_{t=1}^{T} \sum_i \bar{s}_{i\mathrm{LS},t} \Delta \nu_{it}. \tag{68}$$

Matched households have different fixed preference baskets. Retained changes in factor supplies and production weights can move relative prices even under common reference growth, while factor-income changes alter households' shares of total spending. Homotheticity therefore does not imply constant aggregate sectoral shares when households differ.

Figure 55: Welfare differences across initial household groups

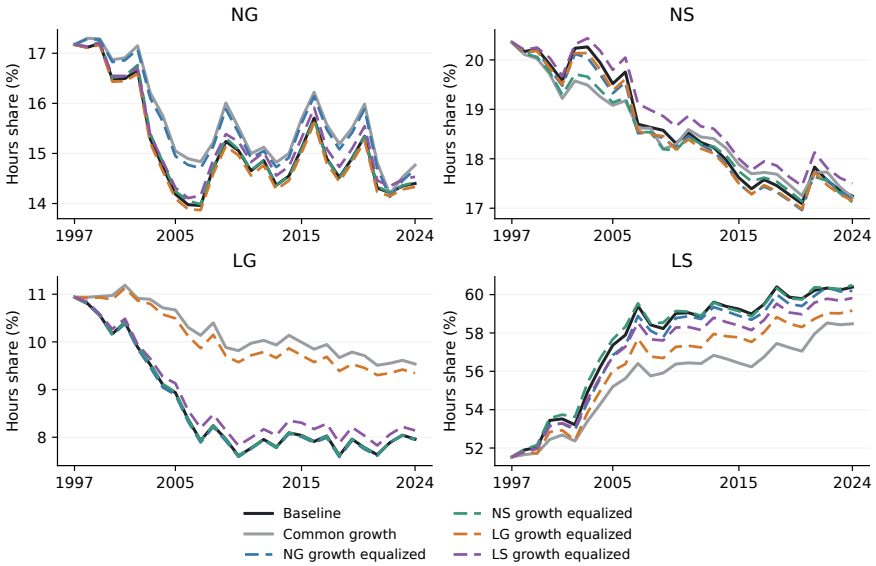


Notes: 2024 baseline-minus-counterfactual welfare under common initial preferences and prices, in percent of initial budgets. Group counts are original CEX households; * marks fewer than 30. Initial sample: Urban households with reference person aged at least 25. Age, sex, marital status and education describe the reference person. Open circles equalize expenditure-decile composition; absent when support is incomplete. Intervals are weighted joint-replica 10th–90th percentiles. Zero-support groups remain explicit.

E.6 Household financing robustness

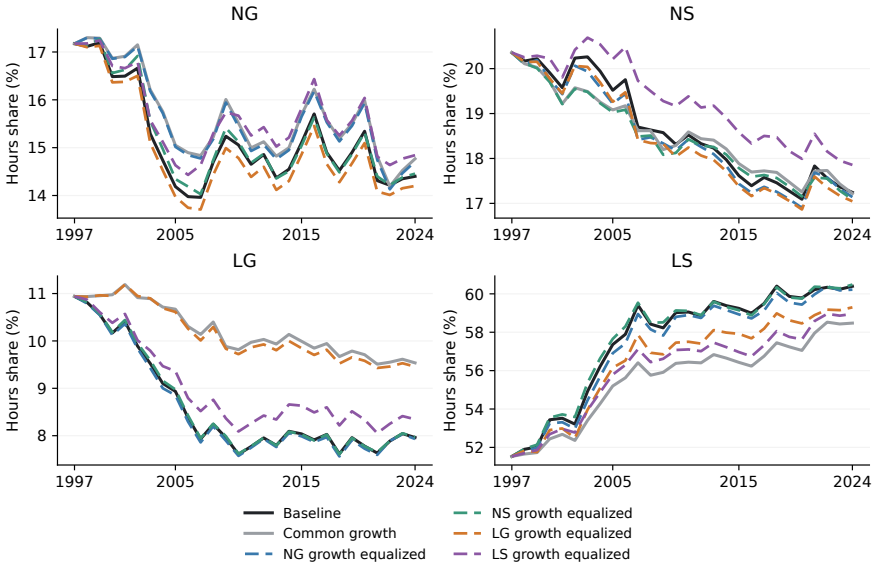
Financing determines how factor-income changes translate into household spending, so it is a substantive modeling assumption. The first comparison preserves every initial budget and the full household population. Let F_{it} be gross factor income and $\bar{F}_t = \sum_i \pi_i F_{it}$. The alternative

Figure 56: Sectoral hours after equalizing one sector's productivity growth



Notes: Hours shares, 1997–2024. The named sector receives the annual four-sector observed-VA-share-weighted mean reference-productivity growth; other sectors retain historical growth. Factor supplies, initial portfolios, and calibrated production-weight paths remain fixed. Black: Baseline. Gray: Common growth in all sectors.

Figure 57: Sectoral hours with average reference-productivity growth preserved



Notes: The named sector receives common reference growth. The same additive log-growth adjustment is applied to each of the other three sectors to preserve the baseline weighted mean in each year. Weights are observed midpoint nominal-VA shares. Black: Baseline. Gray: Common growth in all four sectors. Purchasing power and equilibrium factor returns can change.

rule is

$$y_{it} = \begin{cases} \lambda_t y_{i0} F_{it}/F_{i0}, & F_{i0} > 0, \\ y_{i0} \bar{F}_t/\bar{F}_0, & F_{i0} = 0, \end{cases} \qquad \sum_i \pi_i y_{it} = \bar{F}_t. \tag{69}$$

The common multiplier λ_t balances aggregate budgets, with $\lambda_0 = 1$. Positive-resource households' spending follows their own factor-income growth, subject to this adjustment. Households

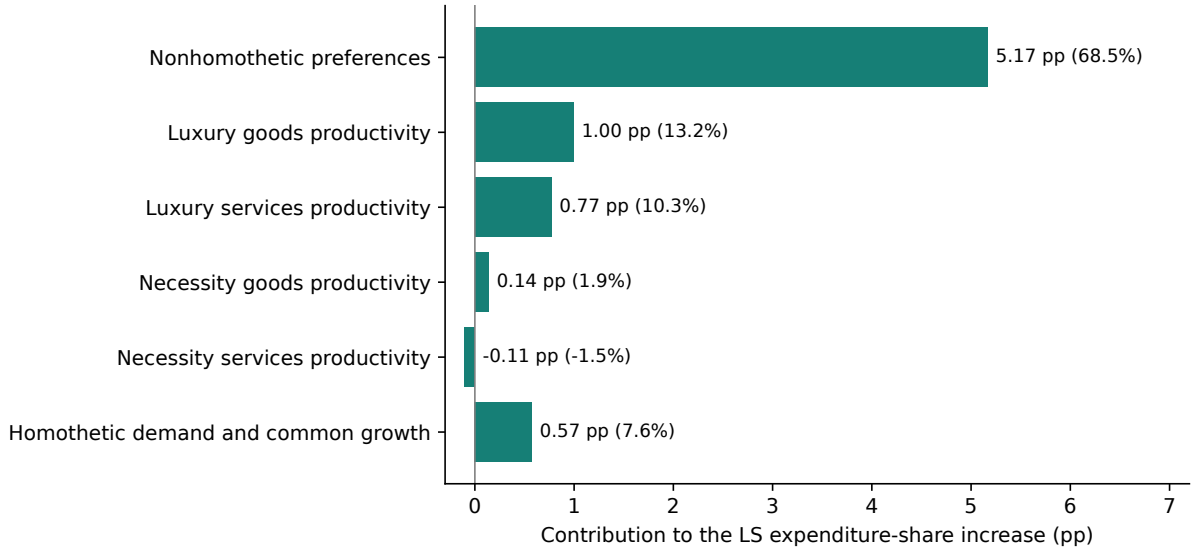
Figure 58: LS hours expansion when luxury sectors retain their distinctive growth



Notes: Changes in the LS share of classified actual hours, 1997–2024; each series subtracts its own 1997 share. Data: KLEMS private industries, fractional four-sector classification, excluding unclassified hours. Model: Five occupation types, unit substitution, and common annual demand shifts. A retained sector follows historical reference-productivity growth; the other sectors receive their renormalized observed-midpoint-VA-weighted mean growth. For LG plus LS, only NG and NS are averaged. All model paths preserve baseline mean reference growth, while household purchasing power can change.

without initial factor income retain their initial share of aggregate income. I solve all four economies again, keeping tastes, initial allocations, assignments, capital holdings, and production and supply paths unchanged. Figure 60 compares the resulting incidence with baseline financing.

Figure 59: Sources of luxury-service expenditure-share expansion



Notes: Shapley contributions, 1997–2024, averaging all 120 orders of the nonhomotheticity and four productivity switches. Remaining expansion under homothetic demand and common growth is shown separately.

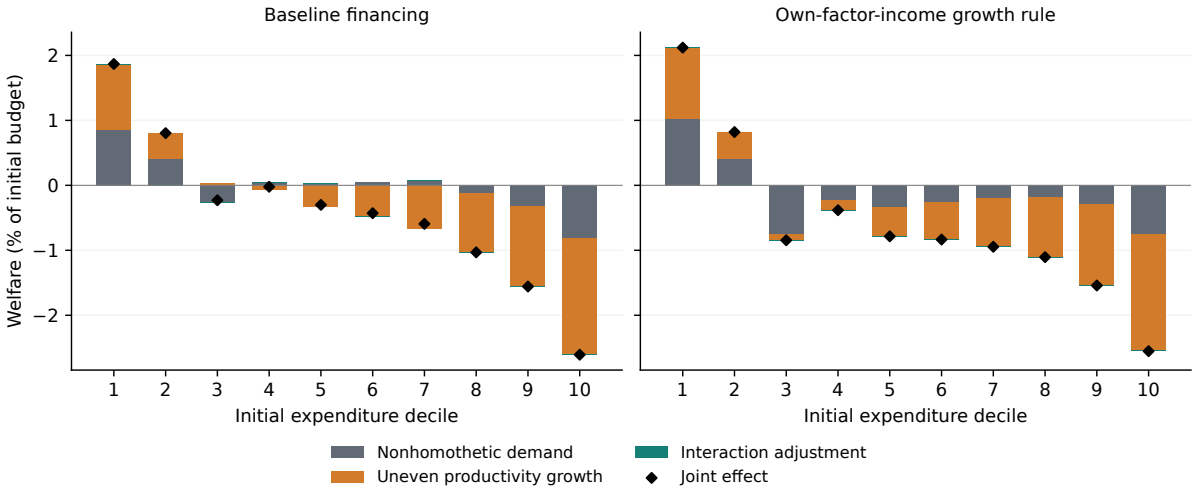
Table 19: Shapley contributions across switching orders

Component	Mean	Minimum	Maximum
NG	0.14	0.11	0.17
NS	-0.11	-0.13	-0.09
LG	1.00	0.86	1.13
LS	0.77	0.62	0.93
Nonhomothetic demand	5.17	4.85	5.49

Notes: Contributions to the 1997–2024 LS expenditure-share increase, in percentage points. Mean, minimum, and maximum over all 120 orders. All configurations use the same replacement-growth path and fixed objects. The ranges are not confidence intervals.

The alternative preserves gains at the bottom and losses at the top of the expenditure distribution, but lowers the gains from nonhomothetic demand in the middle deciles, turning some small gains into losses. Across all three comparisons, the largest revision to a decile mean is 0.62 percentage points. Thus, financing matters for the precise incidence within the distribution even though the broad contrast between the bottom and top survives.

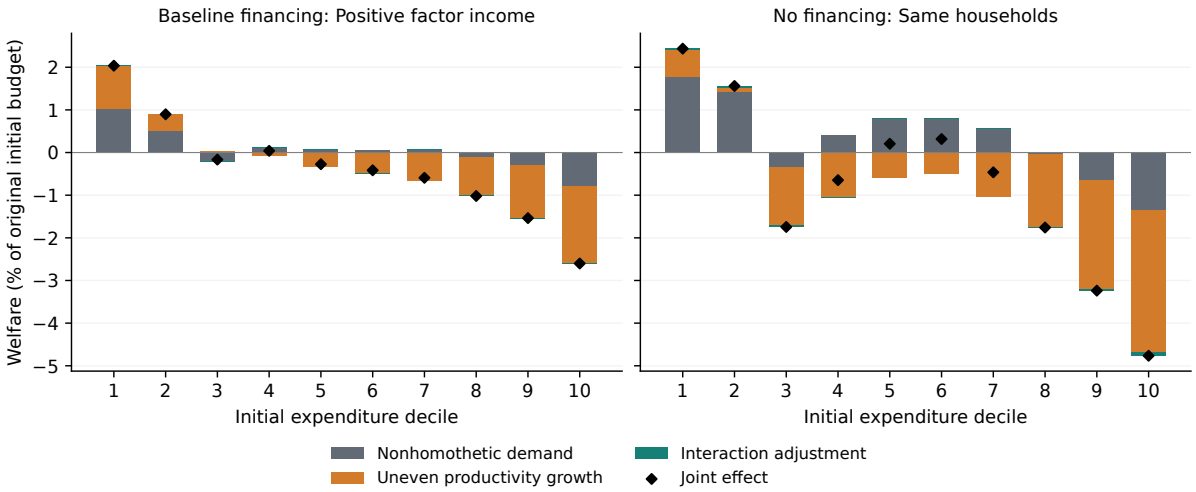
Figure 60: Welfare under two balanced financing rules



Notes: 2024 baseline-minus-counterfactual welfare under each financing rule. Bars show nonhomothetic demand, uneven productivity growth, and their interaction adjustment; black diamonds show the joint effect. The interaction is joint minus the two separate effects. Full initial population, identical initial budgets, fixed expenditure ranks, common reference preferences and prices. All equilibria are re-solved. The alternative indexes positive-resource household budgets to own factor-income growth and balances them with a common multiplier.

Removing financing altogether is a broader stress test. Households with no factor endowments would have zero consumption, so percentage welfare is undefined for them. Figure 61 compares the same positive-factor-income households on both sides; they represent 98.04% of the original population weight. Reporting uses their original expenditure-decile assignments, reference preferences and prices, and original after-financing budget denominators. The no-financing initial equilibrium is solved again and the homothetic benchmark is matched to that new allocation. Thus, this experiment also changes the initial spending fit and prices; it does not isolate only the subsequent transmission rule. No household is removed from the baseline or the balanced-financing comparison.

Figure 61: Welfare with and without financing: A positive-resource stress test



Notes: Same positive-factor-income households under both rules. Bars show nonhomothetic demand, uneven productivity growth, and their interaction adjustment; black diamonds show the joint effect. Original baseline expenditure ranks and initial total-budget denominators, with common initial baseline nonhomothetic preferences and prices for valuation. The no-financing initial allocation changes and its homothetic match is reanchored. Means condition within the original decile on eligible households.

E.7 Occupation aggregation and substitution

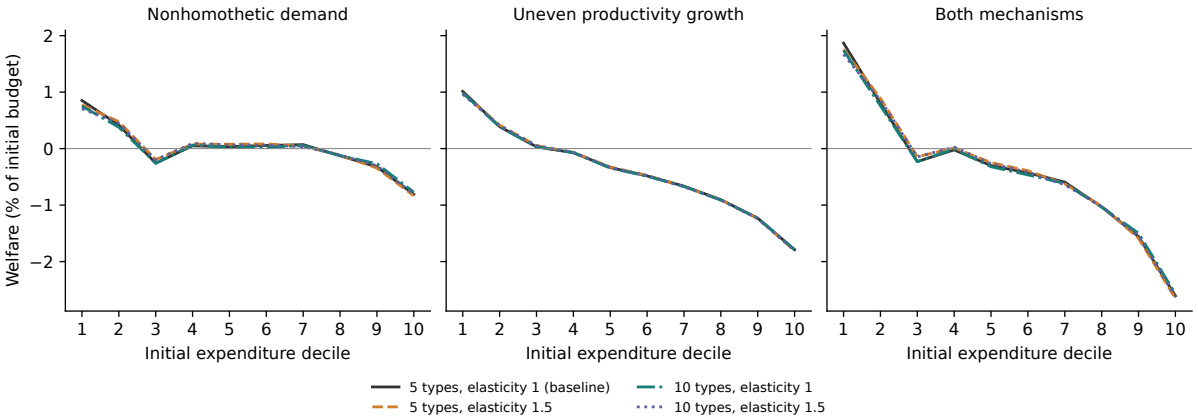
Figures 62 and 63 compare all four combinations of five or ten occupation groups and labor-substitution elasticities of one or 1.5. Every version uses common annual type-demand shifts, the corrected 2003 bridge, and the reconciled aggregate KLEMS labor-services target. The occupation assignments are rebuilt at the relevant aggregation, retaining parent household masses, total hours, capital holdings, expenditure ranks, and initial budgets. Within each calibration, common type shifts match its relative wage targets and remain fixed across counterfactuals.

Relative to the five-type, unit-elasticity baseline, the maximum change in decile-average welfare is 0.11 percentage points with five types and elasticity 1.5, 0.13 with ten types and elasticity one, and 0.20 with ten types and elasticity 1.5. The occupational U-shape in gains from nonhomothetic demand survives, while higher substitution moderates the gains at the ends of the type distribution. The five- and ten-type panels represent different aggregations, so individual type labels do not identify the same group across them. These comparisons support the stability of the main incidence pattern within this set of specifications. They do not identify the production elasticity or remove uncertainty about imputed joint portfolios and financing.

E.8 Fixed occupation-demand weights

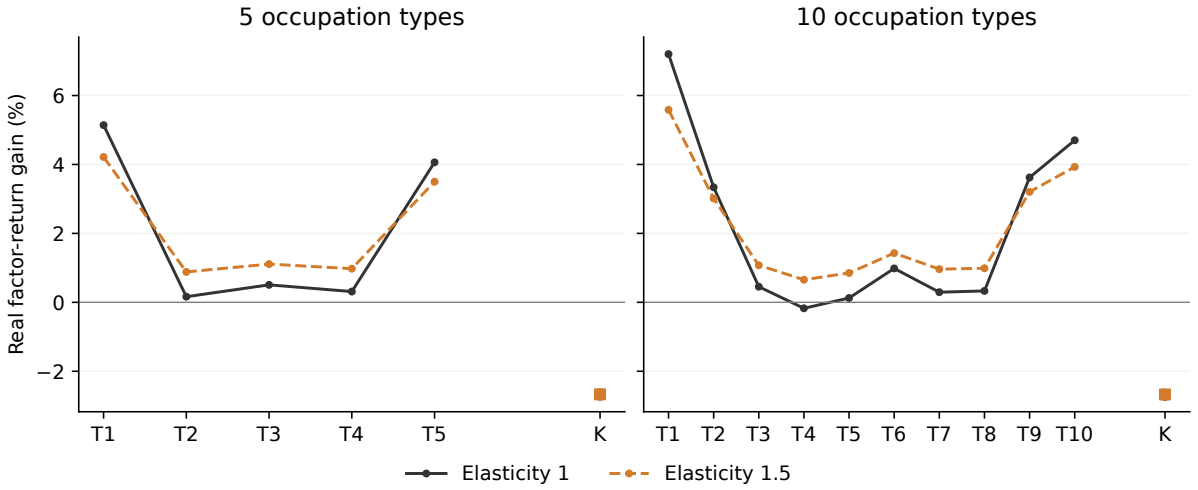
The final comparison asks whether fitting historical relative wage growth drives the counterfactual incidence. I set $g_{jt} = 0$ for every type and year, fixing $a_{sjt} = a_{sj0}$. Relative wages then adjust in equilibrium without matching the ORG growth targets. Initial allocations, household portfolios, capital holdings, tastes, factor quantities, common labor quality and capital-share paths remain unchanged. I retain the historical reference-return convention and recompute the productivity correction for constant type weights. This removes the type-weight contribution

Figure 62: Welfare under alternative occupation counts and substitution elasticities



Notes: 2024 baseline-minus-counterfactual welfare within each calibration, in percent of initial household budgets. All four variants use common annual type-demand shifts. Fixed parent expenditure ranks and common valuation definitions. Baseline: Five types and unit substitution.

Figure 63: Factor-return gains under alternative occupation specifications



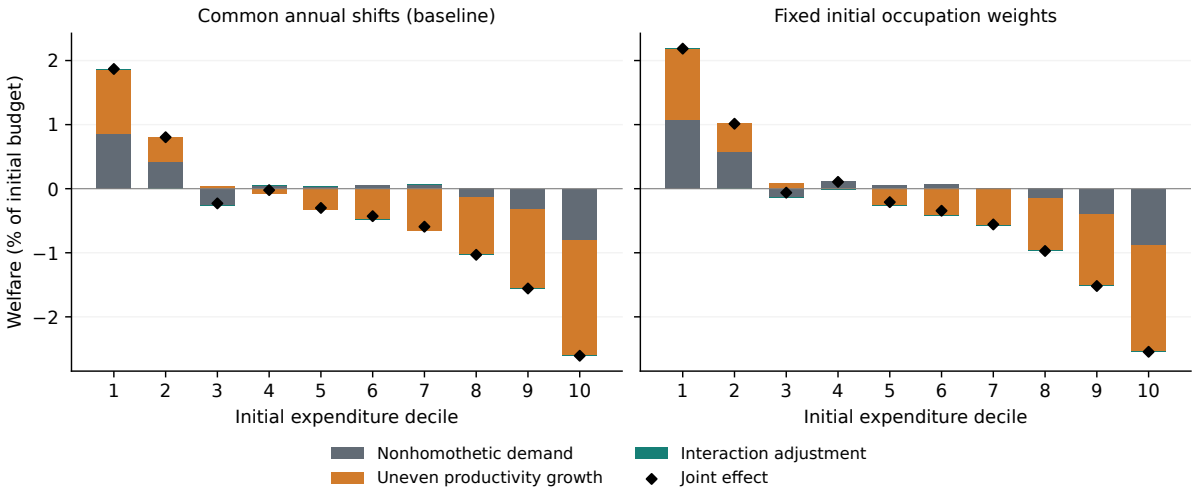
Notes: Real factor-return gains from nonhomothetic demand, 100 × (baseline/homothetic − 1), in 2024. Five-type and ten-type panels share a vertical scale; types are ordered by initial occupation pay within each partition. K denotes the capital rental rate. Each economy uses its own Divisia deflator.

while preserving the interpretation of the KLEMS reference-productivity input; no residual productivity path is fitted to restore wage fit.

Figure 64 compares welfare under annual common shifts and fixed initial weights. Figure 65 shows the cost in relative-wage fit. Tables 20 and 21 report the corresponding LS expansion and real factor-return gains. This exercise changes the demand-weight restriction without changing occupation count or the substitution elasticity.

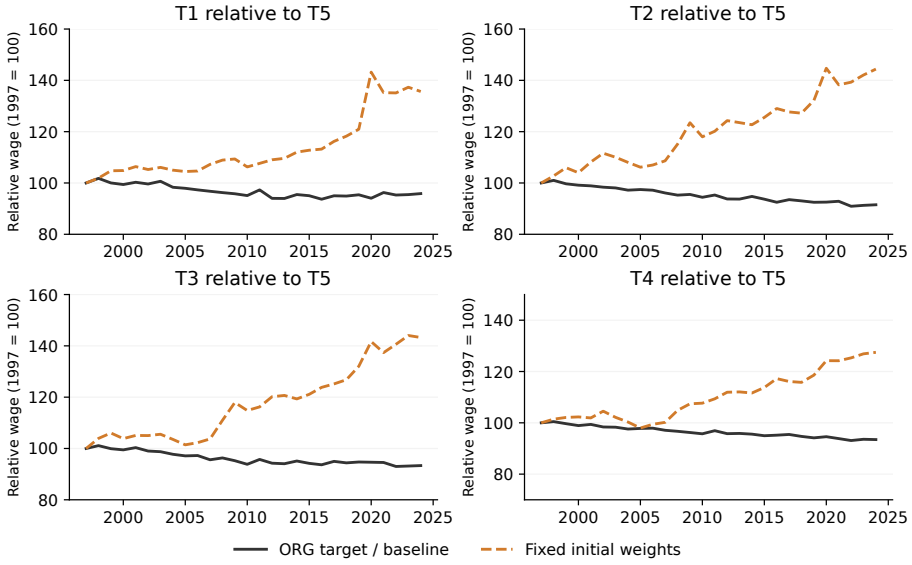
The U-shape in occupational gains from nonhomothetic demand survives, as do higher welfare at the bottom and lower welfare at the top of the expenditure distribution. The largest change in a decile mean is 0.32 percentage points across the three comparisons. The fixed-weight model fits relative wage growth substantially less well: The root-mean-square error across the four

Figure 64: Welfare with annual common shifts and fixed occupation weights



Notes: Five occupation types and unit substitution throughout. Bars show nonhomothetic demand, uneven productivity growth, and their interaction adjustment; black diamonds show the joint effect. Fixed-weight version: All four independent common log demand shifts set to zero. All four economies are solved again, preserving initial household budgets, portfolios, expenditure ranks, and the common welfare valuation.

Figure 65: Relative-wage fit lost with fixed occupation weights



Notes: Hourly wages relative to T5, rebased to 100 in 1997. The annual-shift baseline matches the four adjusted ORG targets. The fixed-weight model imposes no later relative-wage targets. Supplies, common labor quality, and the initial occupation weights are unchanged. All panels use equal vertical spans and tick spacing.

target ratios over 1998–2024 is 22.35 log points. Hence, annual shifts matter for the historical wage fit, but fitting these paths does not generate the main counterfactual incidence pattern.

The sectoral productivity comparison is also stable when occupation weights are fixed. Figure 66 repeats the mean-preserving experiments under both specifications. Fixing the weights raises baseline LS hours expansion from 8.86 to 9.52 percentage points, but changes the reduction associated with equalizing each sector's growth by less than 0.02 points. The LG reduction is 1.10 points with annual shifts and 1.11 with fixed weights; the LS reduction is 1.44 points in

Table 20: LS expansion with annual common shifts and fixed weights

Economy	VA: Annual	VA: Fixed	Hours: Annual	Hours: Fixed
Baseline	7.54	6.99	8.86	9.52
Matched homothetic	2.39	1.98	3.90	4.71
Common productivity growth	5.76	5.19	6.95	7.59
Both changes	0.57	0.16	1.86	2.66

Notes: 1997–2024 percentage-point changes in the LS share of nominal value added and actual hours. Annual: Common annual type-demand shifts. Fixed: Initial occupation weights throughout. The four rows change preferences, reference-productivity growth, or both within each labor-demand specification.

Table 21: Factor-return gains from nonhomotheticity with annual and fixed weights

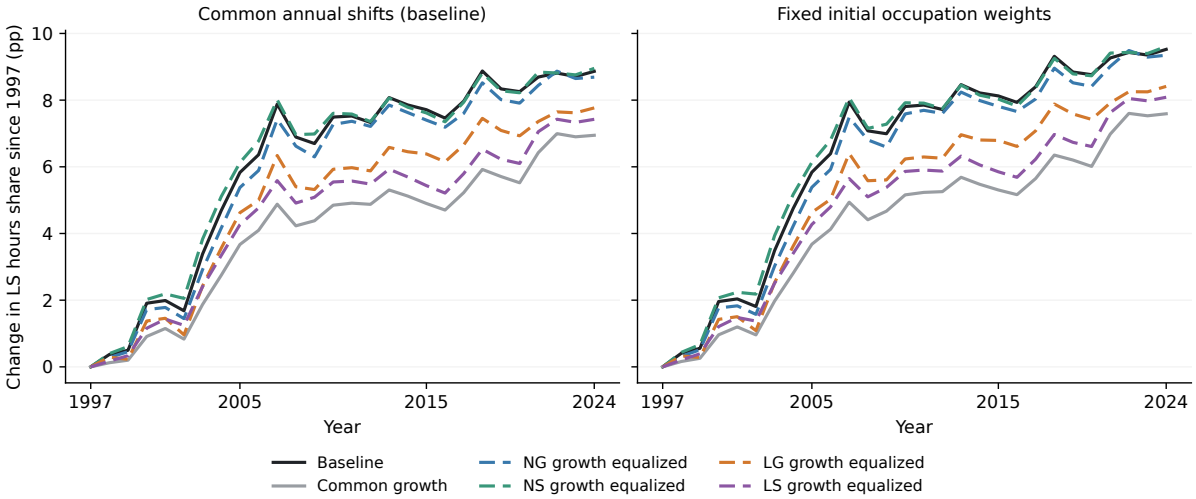
Factor	Annual shifts	Fixed weights
T1	5.14	5.31
T2	0.16	0.44
T3	0.51	0.80
T4	0.31	0.62
T5	4.06	4.24
Capital	-2.68	-2.54

Notes: 2024 real return gains, $100 \times$ (baseline/homothetic $- 1$), using each economy's Divisia deflator. Five types, unit substitution, identical initial portfolios and common quality. Values are factor-return changes, not household welfare changes.

both. Thus, these sectoral comparisons do not depend on fitting the historical relative wage paths.

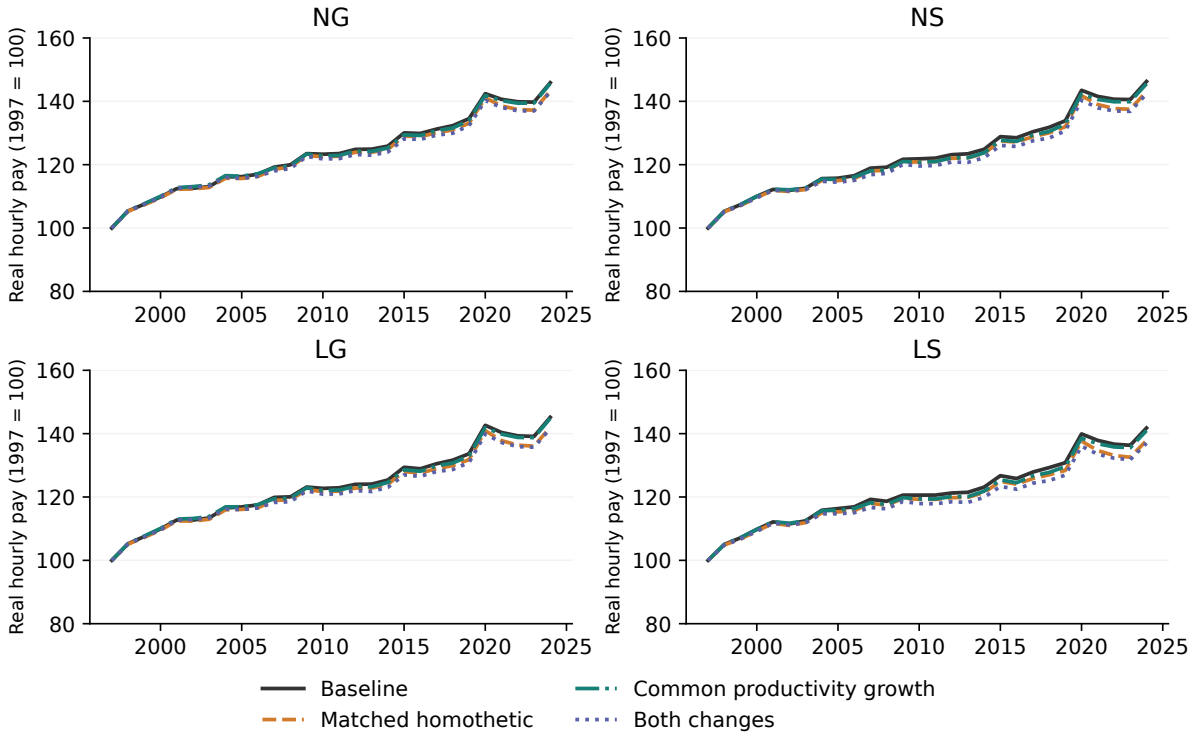
Figure 67 shows sectoral pay in the four fixed-weight economies. Sectoral pay equals the labor bill divided by actual hours. Thus, these averages change with both occupation wages and

Figure 66: Sectoral productivity experiments with annual shifts and fixed occupation weights



Notes: Changes in the LS hours share from 1997, in percentage points. Each panel compares its own baseline, common reference growth, and four experiments equalizing the named sector's growth while offsetting growth elsewhere to preserve the annual observed-VA-weighted mean. Left: Common annual occupation shifts. Right: Fixed initial occupation weights, with the reference-productivity correction recomputed and no relative-wage refit. Factor quantities, labor quality, capital shares, preferences, and initial household portfolios are unchanged. Both panels share a vertical scale.

Figure 67: Counterfactual sectoral real hourly pay with fixed occupation weights



Notes: Sectoral labor bills divided by actual hours, deflated by each economy's aggregate Divisia price index and rebased to 100 in 1997. Initial occupation-demand weights remain fixed in all four economies. The comparisons change preferences, the reference-productivity growth path, or both; factor supplies, common quality, capital shares, and initial portfolios remain fixed. All four panels have equal vertical spans and tick spacing.

the mix of workers employed in each sector. They complement the fixed-type relative-wage fit comparison in Figure 65.