Econometrica Supplementary Material

SUPPLEMENT TO "THE RACE BETWEEN PREFERENCES AND TECHNOLOGY"

(*Econometrica*, Vol. 91, No. 1, January 2023, 227–261)

JOACHIM HUBMER
Department of Economics, University of Pennsylvania

APPENDIX A: DATA

A.1. *Input–Output Tables*

THE BEA'S I–O TABLES, available quinquennially 1982–2012, allow to construct a panel dataset of good-level labor shares. Below I describe the necessary data construction steps and relate the resulting aggregate labor share series to a benchmark from the BLS.

Proprietors' Income and Taxes. I follow Gollin (2002) and Valentinyi and Herrendorf (2008) in employing the economy-wide assumption for splitting the ambiguous part of proprietors' income into payments to labor and capital. To do so requires using the BEA's GDP-by-Industry tables, available only at the 2-digit NAICS level. In both data sets, industry value added is broken down in three parts: compensation of employees (*coe*); production taxes and subsidies (*tx*); and a residual called gross operating surplus (*gos*). In the GDP-by-Industry tables, *gos* can be further decomposed into a part that is unambiguous capital income (corporate gross operating surplus plus noncorporate consumption of fixed capital), and into ambiguous income (noncorporate net operating surplus). Then I compute $prop_j$, for each 2-digit industry j, defined as the fraction of GOS that is ambiguous income. Subsequently, I map detailed I–O industries i to 2-digit NAICS industries j using the official concordance, defining $j(i)$. Finally, I reallocate a portion

$$prop_{j(i)} \times gos_i \times \frac{coe_i}{coe_i + (1 - prop_{j(i)}) \times gos_i}$$

from capital income gos_i to labor income coe_i. I exclude taxes when calculating labor shares: labor shares are computed as $\frac{coe_i}{coe_i + gos_i}$ after reallocating part of proprietors' income.

From Industry- to Good-Level Labor Shares. For expositional simplicity, assume that each good i is produced only by a single industry i, and each industry produces a single good.[1] Let $N_t = |I_t|$ denote the number of goods (industries) in year t. Define $\beta_{it} \in (0, 1]$ as the ratio of value added to gross output in industry i, and $\Gamma_{ijt} \in [0, 1]$ as the good j cost share in the intermediate input bundle used for production of good i (i.e., $\sum_{j \in I_t} \Gamma_{ijt} = 1$). Then the overall labor share of good i, θ^L_{it}, solves the following linear system:

$$\theta^L_{it} = \beta_{it} \tilde{\theta}^L_{it} + (1 - \beta_{it}) \sum_{j \in I_t} \Gamma_{ijt} \theta^L_{jt}, \quad \text{for } i \in I_t. \tag{A.1}$$

Joachim Hubmer: jhubmer@sas.upenn.edu

[1]In my treatment of the data, I account for the fact that a single good may be produced by multiple industries, and vice versa.

 https://doi.org/10.3982/ECTA18580

Define the $N_t \times N_t$ matrix $\Gamma_t = [\Gamma_{ijt}]_{i \in I_t, j \in I_t}$, let $D(z)$ denote the diagonal matrix corresponding to a vector z, and let E_n denote the identity matrix of size n. For all other objects defined earlier, z_t denotes the vector $(z_{it})_{i \in I_t}$.

The matrix version of (A.1) is

$$\theta_t^L = D(\beta_t)\tilde{\theta}_t^L + D(\vec{1} - \beta_t)\Gamma_t \theta_t^L,$$

from which we can solve for the vector of final good labor shares,[2]

$$\theta_t^L = \left[E_{N_t} - D(\vec{1} - \beta_t)\Gamma_t\right]^{-1} D(\beta_t)\tilde{\theta}_t^L.$$

Industry Classifications. I–O industry classifications are time-varying, based on SIC industries up to 1992, and on NAICS industries after 1997. The changes from 1982–1992, as well as 1997–2012, are minor, and I create manual concordances for these years based on the I–O Tables' documentation. For the link between 1992 and 1997, I combine three official concordances: I–O 1992 to SIC, SIC to NAICS, and NAICS to I–O 1997. All concordances are weighted by final demand expenditure shares.

Labor Shares of Aggregate Final Demand Components. Given final good labor shares θ_{it}^L, for $i = 1, \ldots, I_t$, $t = 1982, \ldots, 2012$, and final demand component expenditure weights ω_{it}^f for $f \in \{\text{PCE}, \text{PFI}, \text{GP}, \text{NX}\}$, I compute the labor share of component f as $\bar{\theta}_t^{L,f} = \sum_{i \in I_t} \omega_{it}^f \theta_{it}^L$. Figure A.1 reports labor shares by year and final demand component in the left panel, and aggregating various components in the right panel.[3] The right panel shows

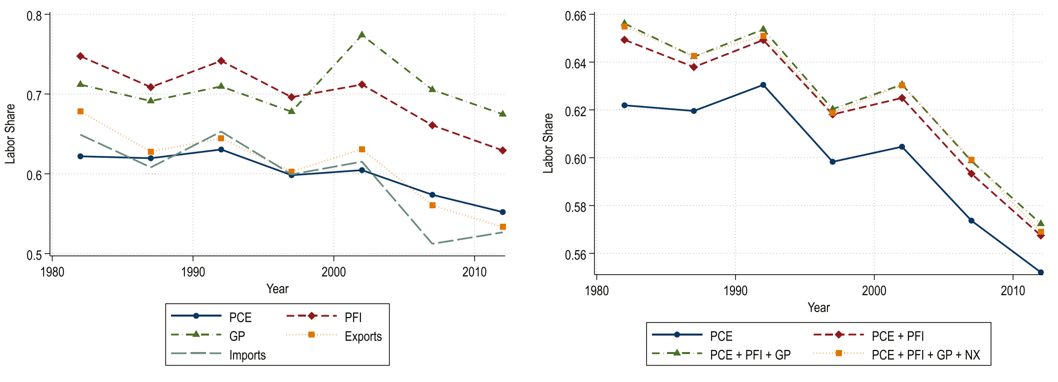

(a) Final demand component labor shares (b) Aggregate labor share

FIGURE A.1.—Labor shares of final demand components. Source: BEA I–O Tables. Aggregate final demand is the sum of personal consumption expenditures (PCE), private investment (PFI), government purchases (GP), and net exports (NX).

[2]The matrix $A_t^{-1} \equiv [E_{N_t} - D(\vec{1} - \beta_t)\Gamma_t]^{-1}$ is the Leontief inverse. Formally, since A_t is an M-matrix, all entries of A_t^{-1} are nonnegative. Moreover, it can be directly verified that the rows of $A_t^{-1}D(\beta_t)$ sum to one, such that indeed good-level labor shares are weighted averages of industry-level labor shares: The claim is that $A_t^{-1}D(\beta_t)\vec{1} = \vec{1}$. This is true if and only if $\vec{1} = D(\beta_t)^{-1}A_t\vec{1}$. Since the rows of Γ_t sum to one by assumption (i.e., $\Gamma_t\vec{1} = \vec{1}$), we have that $D(\beta_t)^{-1}A_t\vec{1} = D(\beta_t)^{-1}[E_{N_t} - D(\vec{1} - \beta_t)\Gamma_t]\vec{1} = D(\beta_t)^{-1}\vec{1} - D(\beta_t)^{-1}D(\vec{1} - \beta_t)\vec{1} = D(\beta_t)^{-1}(E_{N_t} - D(\vec{1} - \beta_t))\vec{1} = D(\beta_t)^{-1}(D(\beta_t))\vec{1} = \vec{1}$, which proves the claim.

[3]For imports, this correspond to the (hypothetical) labor share of the basket of goods imported into the U.S., using the same technology that is currently used to produce domestic output in these sectors.

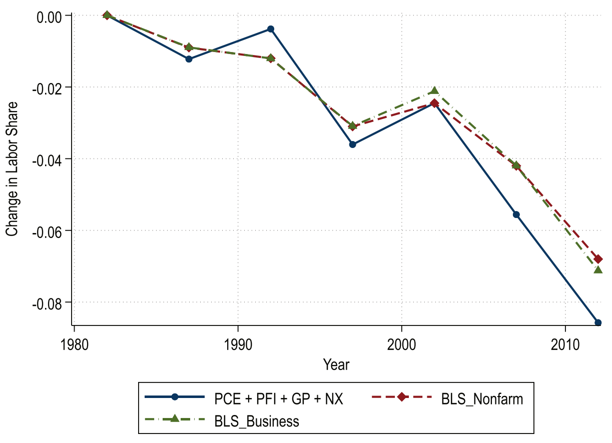

FIGURE A.2.—Comparison of I–O Table labor shares to BLS series. Source: BEA I–O Tables, BLS. BLS_Nonfarm refers to the BLS series for the nonfarm business sector, BLS_Business to the BLS series for the business sector. PCE + PFI + GP + NX refers to the series constructed using I–O Table labor shares for each good and service, weighted by total final demand shares.

that adding investment to consumption spending implies an about one percentage point steeper aggregate labor share decline. On the other hand, the remaining final demand components—government expenditure and net exports—do not meaningfully impact the aggregate labor share evolution.

Treatment of Imports. The baseline approach in this paper assumes a closed economy, with final demand consisting of consumption and investment, both produced fully domestically. Figure A.1(b) shows that excluding net exports (and government purchases) does not affect the aggregate labor share materially. To further rule out confounding effects on the cross-sectional findings, robustness checks consider an alternative definition of the good-level labor share that decomposes a dollar of expenditure on good i in period t into payments to domestic labor, to domestic capital, and to imported value added (Table IV for production; Section B.1 for consumption). Specifically, the I–O Tables contain the total value of imports, for each good (but not by using industry or final demand component). Thus, it is necessary to make a simplifying assumption; following the BEA, I assume that each industry uses imports of any good in the same proportion as the imports-to-domestic supply ratio of that good. Given this assumption, one can solve for the domestic total requirements matrix, and subsequently for the import share and domestic labor share of each good (in each period). See Horowitz and Planting (2014) for details.

Comparison to BLS Labor Shares. Figure A.2 compares the self-constructed labor share series, based on the I–O Tables, to the series provided by the BLS. Conceptually, the BLS series for the business sector (dashed line with triangles) corresponds closely to the self-constructed series for total final demand (PCE + PFI + GP + NX). Reassuringly, the series largely agree.

A.2. *Consumer Expenditure Survey (CEX)*

Sample Selection. The sample is restricted to households with household heads aged 25–65, a full-year of interview coverage (four quarterly interviews), and complete income responses. The latter concept is captured by the variable *RESPSTAT*.

Owner-Occupied Housing. In general, the CEX measures out-of-pocket expenditures. Hence, while it reports the appropriate rent for renters, for homeowners it reports cash expenditures associated with owning a house (mortgage interest, property taxes, home insurance, maintenance expenses, etc.). Fortunately, the rental equivalence of owning a house is recorded as well. I treat this equivalent rent both as a component of consumption and of income. To avoid double counting, out-of-pocket expenditures for homeowners have to be subtracted. Since homeowners' rental equivalence is not reported prior to 1982, I impute it based on later survey waves (1982–1990).[4]

Income Concept. I use after-tax household income as reported in the CEX (*FINCATAX*) and add the net rental equivalence of homeowners (as explained above).

Diary Survey Items. I mainly rely on the CEX interview survey, which covers the majority of household expenditures. The interview survey is missing expenses on housekeeping supplies, personal care products, and nonprescription drugs, which amount to 5–15% of total expenditures and which are reported in the diary survey. Consequently, I impute missing expenditures based on diary survey data.[5]

Treatment of Zeros. For some goods with positive aggregate CEX expenditures in a given year, there are households with zero recorded expenditure. This could be the case if either the households forgot to record the item, or they simply did not spend anything on it. I impose a lower bound on household expenditure shares equal to one-tenth of a good's aggregate CEX expenditure share in that year in order to be able to take the logarithm and not have to drop these households.

A.3. *Linking CEX and I–O Tables*

The link from CEX data to the BEA's Detailed I–O Tables is based on a manual concordance by Levinson and O'Brien (2019). This concordance only covers the interview survey. I add diary survey items manually, as well as a few interview survey items that are not part of their concordance (rental equivalence of homeowners, used car expenses).

Producer versus Purchaser Prices. Expenditures in the CEX are denominated in purchaser prices, whereas the I–O Tables are in producer prices. The difference between the former and the latter is a set of margins (wholesale, retail, and transportation). The I–O Tables contain the necessary information to convert expenditures in purchaser prices to expenditures in producer prices. In particular, the Use Table contains for each dollar of final demand expenditure on a good $i \in I$: the fraction of that dollar that is recorded as revenue by the producer of good i, as well as the fractions going to the wholesale sector, retail sector, and various transportation sectors. I reallocate CEX spending according to this map, so that consumer demand is specified in producer prices (production data is already in producer prices).

[4]Specifically, I predict it by regressing the expenditure share of owner occupied housing on log income, log total other expenditures, and demographic controls (reference person's age, race, and sex; household size; region; number of earners).

[5]For each of these consumption categories, I predict their expenditure share by regressing annual household expenditure as a fraction of household income on log income, demographic controls (see above), and calendar year.

Aggregate Consumption in the CEX versus PCE. A concern with the CEX data is that its representativeness of aggregate consumption (PCE) has been declining (see Garner, Janini, Passero, Paszkiewicz, and Vendemia (2006)). First, my estimated income elasticities are robust to measurement error as long as it is of the form discussed in Aguiar and Bils (2015); that is, as long as measurement error is household-specific and/or good-specific, and not household-good-specific. Second, these discrepancies raise the question of whether it is more appropriate to use aggregate CEX expenditure weights ω_{it}^{CEX} or PCE-based weights ω_{it}^{PCE} for the model analysis. On the one hand, PCE are preferable because they are more reliable for aggregate trends. On the other hand, the estimated income elasticities only correspond to the fraction of consumption expenditures that is recorded in the CEX. I use CEX weights for the baseline model with consumption only (Section 6.1), and PCE weights in Section 6.2 (where I also include aggregate investment demand). Reassuringly, the comparison between model and data is similar in both cases.

A.4. *Equipment Capital Intensities*

Capital Classification. Throughout this paper, I consider a two-way split of capital income into private equipment and software, as well as private structures. This partition corresponds to the one prior to the 14th comprehensive revision of NIPA in 2013, which capitalized a larger set of intellectual property products (IPPs), and classified them as a separate asset category that also includes software. Figure A.3 shows that according to the integrated BEA/BLS estimates, those IPP categories that are excluded from the analysis in this paper—R&D as well as artistic originals—contribute less than 1pp of the roughly 7pp increase in the aggregate capital share.

User Cost of Capital. Mapping nominal capital stocks into factor payments requires an assumption on the required rate of return. A naive strategy would assume that returns on equipment and structures are equal, and thus drop out from equation (1). However, depreciation rates are much larger for equipment; moreover, based on historical experience the expected price decline is also larger. Thus, I employ a standard user cost formula. The

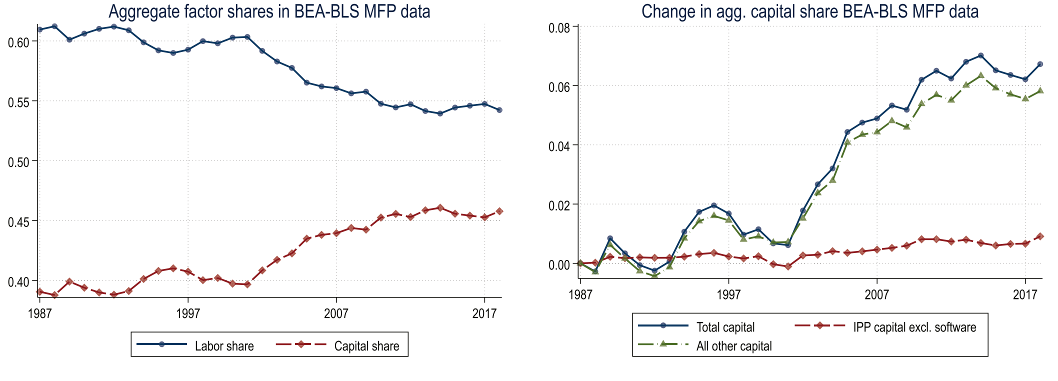

FIGURE A.3.—Decomposition of aggregate capital share in BEA-BLS MFP data. Source: BEA/BLS integrated industry-level production account 1987–2018. In the right panel, "IPP capital excl. software" depicts the change in the capital share due to R&D and artistic originals, while "All other capital" includes equipment, structures, as well as software.

required rate of return on a dollar of capital type $k = E, S$ is given by

$$\tilde{R}_t^k = r_t + \delta_t^k - \left(1 - \delta_t^k\right) \frac{\mathbb{E}_t\left[P_{t+1}^k - P_t^k\right]}{P_t^k}, \tag{A.2}$$

where r_t is a real interest rate, δ_t^k is the depreciation rate of type k capital, and the last term refers to expected price growth. I compute r_t as a weighted average of the cost of debt and equity.[6] To compute δ_t^k by capital type and year, I divide current-cost depreciation by current-cost net stock of capital (FAT Tables 1.1 and 1.3). For the expected price growth term, I use a 5-year moving average of realized price growth. Note that p_t^k refers to the price of type k capital relative to a consumption price index. Specifically, I use the FRED series PERICD for equipment, which refers to the quality-adjusted price of equipment and software, relative to a consumption deflator. For structures, I use the BEA's nonresidential structures (B009RG3Q086SBEA) and residential investment (B011RG3Q086SBEA) deflators, again relative to the same consumption deflator.

Time Variation. In principle, this strategy gives rise to time-varying equipment shares. However, I do not use time-variation in the empirical analysis, since it is quite sensitive to the choice of interest rate (see Figure A.4)—this is because the duration of structures capital is much higher than the one of equipment and software.

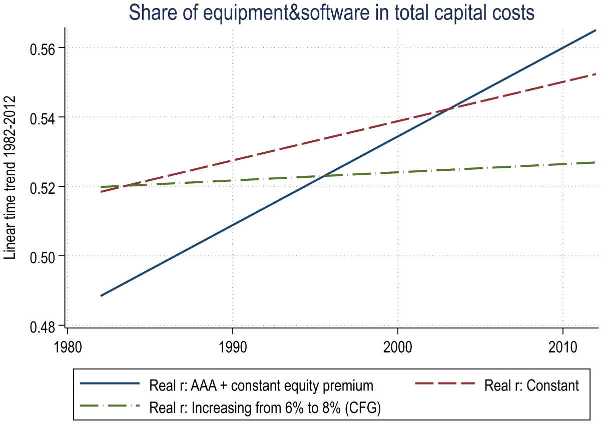

FIGURE A.4.—Time-variation in share of equipment and software in total capital cost (κ_t). The figure reports the aggregate share of equipment and software in total capital costs κ_t, displayed as linear time trend, when using the user cost formula (A.2) and (1), for three different choices of r_t. The solid line corresponds to the realized real interest rate as described in the text (which decreases over time). The dashed line uses the constant time-average of realized returns. The dash-dotted line corresponds to a real rate that linearly increases from 6 to 8%, as in Caballero, Farhi, and Gourinchas (2017).

[6]Specifically, I compute the weights using Table S.5.a of the Integrated Macroeconomic Accounts; computing debt as the sum of line 131 "Debt Securities" and line 135 "Loans," and using line 140 "Equity and investment fund shares" for equity. For the return on debt, I use the AAA bond yield, for equity the 10-year U.S. treasury yield plus a 5% risk premium. I subtract inflation in the form of a 5-year moving average of the CPI from both.

A.5. *Compustat*

Sample Selection. I use data from the Fundamental Annual file, accessed via Wharton Research Data Services, from 1980 to 2014, averaging yearly values over 5-year increments to match the I–O data when applicable (e.g., taking the average 1980–1984 for census year 1982). The sample is limited to firms with Foreign Incorporation Codes (FIC) in the U.S.

Equipment Intensities. These are computed as net equipment and machinery (PPENME) divided by net total property, plant, and equipment (PPENT). I use the time average by industry (converted to goods using the I–O Tables) in a robustness check as alternative measure of $\bar{\kappa}_i$ in column 3 of Table III only.

Markups. For the estimate of μ_{it} in column 5 of Table IV, I rely on the replication of De Loecker, Eeckhout, and Unger (2020) in Hubmer and Restrepo (2021), using the production function approach—also known as ratio estimator—to recover markups as the ratio of the elasticity of gross output (SALE) with respect to variable inputs (COGS) over the revenue share $\frac{\text{COGS}}{\text{SALE}}$.[7]

Industry Classification. I use the NAICS classification in Compustat to aggregate firm-to industry-level data, and then map these into the I–O classification system using the official concordances as described in Appendix A.1.

A.6. *Investment Rates*

Figure A.5 plots the investment (PFI) rate, as a share of PCE + PFI (left panel) as well as as a share of GDP (right panel). In these figures, PFI includes only structures and equipment and software investment, while other IPP investment (R&D and artistic originals) is excluded, for consistency with this paper's approach. The I-Model in Section 6.2 exactly replicates the time series $\frac{\text{PFI}_t}{\text{PFI}_t+\text{PCE}_t}$ depicted in the left panel.

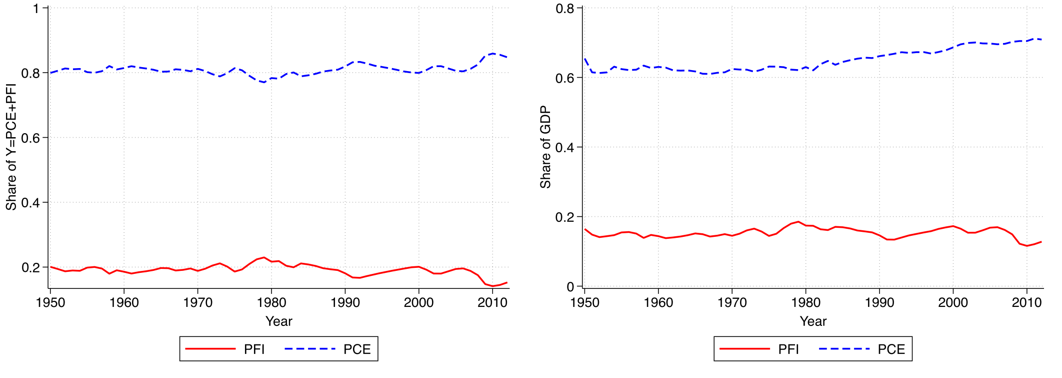

(a) Shares of $Y = C + I$ (b) Shares of GDP

FIGURE A.5.—Nominal investment ratios. Source: BEA. Nominal personal consumption expenditures (PCE) and private fixed investment (PFI).

[7]SALE refers to firm revenue; COGS to cost of goods sold, a measure of variable costs.

APPENDIX B: ADDITIONAL EMPIRICAL RESULTS

B.1. *Consumption and the Labor Share*

Time-Varying Income Percentiles. In Figure 2 in the main text, income percentiles are defined to be constant over time. Figure B.1(a) shows that the relation between household income and household labor shares is hardly changed when using instead time-varying percentiles (such that the level of real income on the horizontal axis is varying over time).

Expenditure Shares by Household Income Quartile. Figure B.1(b) completes Figure 3 in the main text by showing also the second and third quartiles of household income, which are omitted in the main text for clarity. As expected, the relationship is monotone.

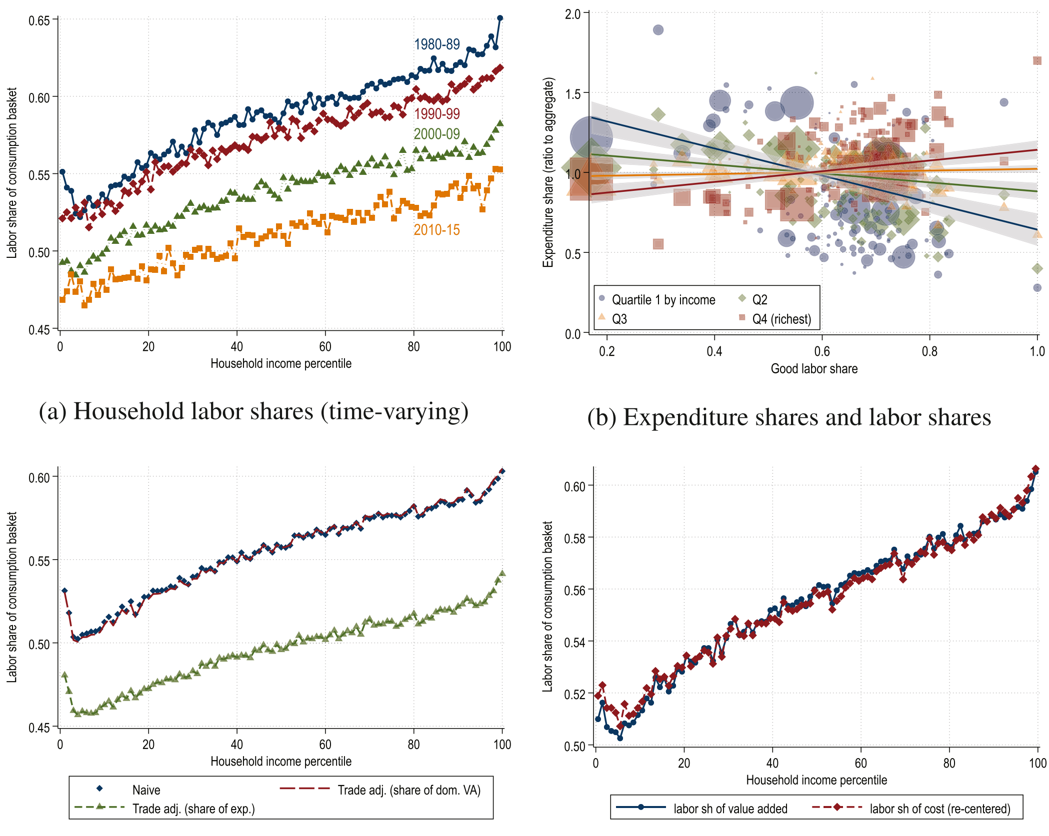

(a) Household labor shares (time-varying)

(b) Expenditure shares and labor shares

(c) Accounting for imported value added

(d) Accounting for markups

FIGURE B.1.—Robustness of relation between household income and labor shares. Source: CEX (household consumption by category and income), BEA I–O Tables (labor and import shares), Compustat (markups). (a) Household labor shares by income percentile, where the latter are time-varying (as opposed to Figure 2 in the main text, where income percentiles are constant). (b) Household expenditure shares computed as average for each household income quartile, displayed as ratio relative to economy-wide aggregate expenditure shares, and ordered by goods' labor shares. (c) Naive refers to baseline labor share measure that assumes all production is domestic. Trade adjustment factors in imports along the value chain, such that a dollar of household spending is decomposed into domestic labor, domestic capital, and imported value added (reported as share of expenditure and of domestic value added). (d) Labor share of cost computed as $\theta_{it}^{L,\text{cost}} = \theta_{it}^{L} \cdot \mu_{it}$ and recentered to baseline labor share of value added.

Accounting for Imports. Figure B.1(c) reports, in addition to the baseline series assuming fully domestic production, two trade-adjusted measures of a household's labor share: the first one is computed as the ratio of a household's spending on domestic labor to its spending on domestic labor and domestic capital, whereas the second one divides by total expenditure instead (which includes spending on imported value added). While the second one results in a lower level, for both measures the cross-sectional difference is virtually the same as in the naive baseline, which assumes that all goods are produced domestically.

Accounting for Markups. In principle, the positive association between a household's income and the labor share of its consumption basket could be driven by (i) labor intensity in production (the baseline interpretation) or by (ii) markups; (ii) would require richer households to disproportionally spend on lower-markup goods. Figure B.1(d) reports, in addition to the baseline measure of the labor share of value added (i + ii), the labor share of cost (i), computed first at the good level as $\theta_{it}^{L,\text{cost}} = \theta_{it}^{L} \cdot \mu_{it}$ (using markups estimated in Compustat, see Appendix A.5) and then aggregated to households. The two series largely agree, confirming that the pattern is indeed driven by differential labor intensity in production.

Sensitivity to Individual Sectors and Weighting. Figure B.2(a) shows that the pattern of higher-income households spending relatively more on labor-intensive goods is not sensitive to the inclusion or exclusion of any particular sector, neither to the weighting scheme: The panel displays the regression coefficients and confidence intervals from regressions of the log expenditure share differential (top-income quartile of households relative to bottom-income quartile) on good-level labor shares. The estimates are positive, statisti-

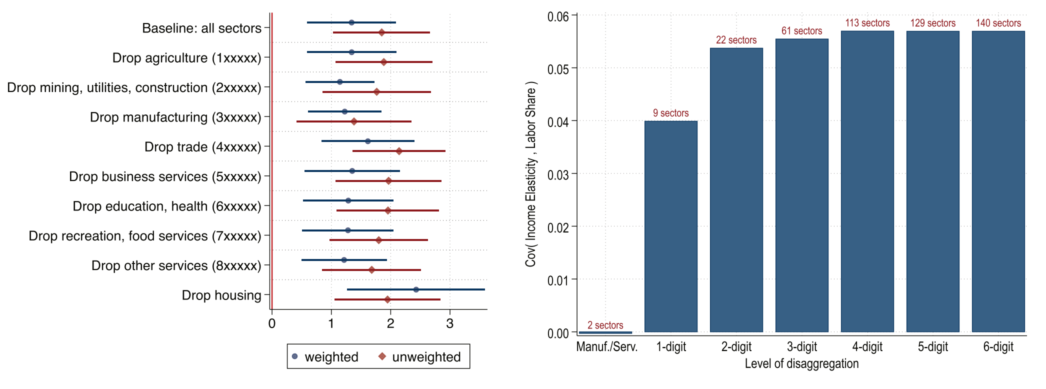

(a) OLS coef. expenditure shares—labor shares (b) Covariance income elasticities—labor shares

FIGURE B.2.—Income elasticity—labor share relation: Robustness and disaggregation. Source: CEX, I–O Tables. (a) Point estimates and 95% confidence intervals from cross-sectional OLS regressions of the expenditure share ratio $\ln(\frac{\omega_i^{Q4}}{\omega_i^{Q1}})$ on good-level labor shares θ_i^L, robust standard errors. Here, $Q4$ ($Q1$) denotes the top (bottom) quartile of households by income. Weighted regressions use aggregate expenditure share weights. Panel (b) displays the (time average over the) cross-sectional covariance of income elasticities and labor shares, for various levels of disaggregation. The rightmost bar corresponds to the level of detail used in the benchmark (Detailed I–O Tables, 6-digit classification, 140 goods in I–O Tables with active link to at least one CEX category). The second from the right corresponds to 5-digit I–O sectors (129 goods), . . . , all the way to 1-digit sectors (9 goods) and manufacturing versus services (2 goods).

cally significant, and stable across specifications. Generally, this paper uses aggregate expenditure share weights. Unweighted, the association is even stronger. A concern would be if any particular sector, such as housing or health care, were to drive this correlation. Reassuringly, when repeating this exercise while systematically dropping all goods within a 1-digit sector, and cycling through all sectors, the pattern is stable.

Covariance by Aggregation Level. Figure B.2(b) displays the cross-sectional covariance between labor shares and income elasticities, for varying levels of disaggregation. Interestingly, when considering only a two-way split of consumption categories into services and manufacturing, the covariance is close to zero. This is because the level difference between the labor share of the manufacturing sector and the one of the services sector is minor (in fact, the manufacturing labor share used to be higher than the one of services; this pattern reversed over time). The 1-digit level (9 goods) already captures almost two-thirds of the covariance; the 2-digit (22 goods) level captures more than 90% of the variation.

Income Elasticities and Labor Shares. Figure B.3 displays the estimated income elasticities, including 95% confidence intervals, as a function of goods' labor shares.

Summary Statistics at 2-Digit Level. Table B.I summarizes expenditure shares, raw expenditure share differences between high- and low-income households, estimated income elasticities, as well as labor shares. The reported values are aggregated to the 2-digit I–O level, and averaged over time. For example, I–O code 3122A0 Tobacco Product Manufacturing is subsumed in I–O Sector 31. Observe that the raw expenditure share differential by income group closely aligns with the pattern of income elasticities, reflecting that estimation of the latter is based on variation in the former.

Price and Taste Heterogeneity. Identification of the income elasticities from cross-sectional data requires that unobserved cross-sectional heterogeneity in prices or tastes is

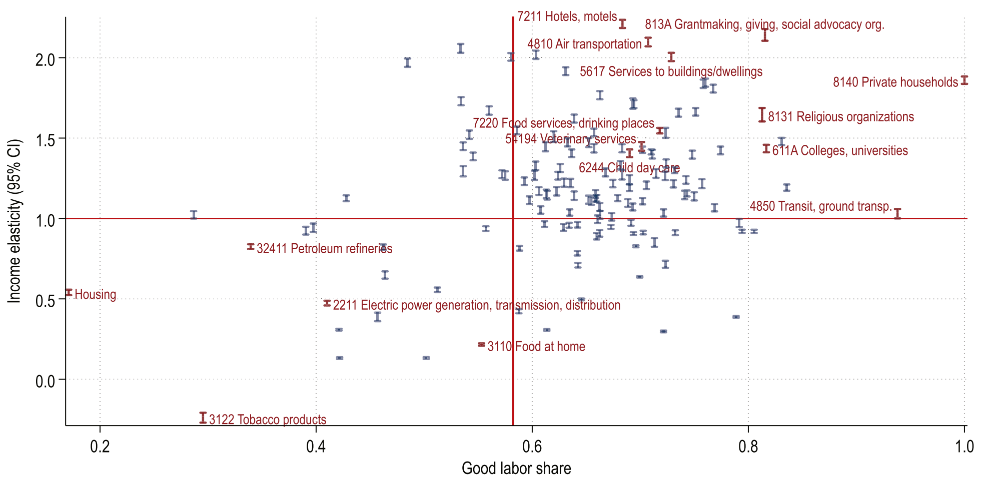

FIGURE B.3.—Estimated income elasticities and labor shares in the cross-section. Source: CEX, I–O Tables. Income elasticities and labor shares averaged over time. Income elasticities reported as 95% confidence intervals, corresponding to pooling regression (10) over all sample years. The vertical line indicates the aggregate labor share. The same examples as in Figure 3 are highlighted.

TABLE B.I

SUMMARY STATISTICS AT THE 2-DIGIT I–O LEVEL.

I–O Sector (2-digit)		ω_i^{CEX}	ω_i^{PCE}	$\ln(\frac{\omega_i^{Q4}}{\omega_i^{Q1}})$	$(\gamma_i - 1)$	θ_i^L	$\theta_i^L - \bar{\theta}^L$
11	Agriculture (nurseries, floriculture)	0.001	0.001	0.692	0.812	0.534	−0.046
21	Mining	0.002	0.000	0.044	0.041	0.291	−0.289
22	Utilities	0.054	0.030	−0.523	−0.299	0.434	−0.146
23	Construction	0.037	0.000	0.987	0.667	0.751	0.171
31	Manufacturing I (food, apparel...)	0.112	0.098	−0.452	−0.559	0.564	−0.016
32	Manufacturing II (wood, chemical...)	0.031	0.049	−0.282	−0.210	0.404	−0.176
33	Manufacturing III (cars, electronics...)	0.123	0.058	0.305	0.170	0.661	0.081
42	Wholesale trade	0.053	0.036	−0.211	−0.353	0.699	0.119
48	Transportation	0.024	0.025	0.408	0.391	0.725	0.145
49	Warehousing, couriers, postal service	0.000	0.001	−0.150	−0.466	0.765	0.184
4A	Retail trade	0.114	0.108	−0.122	−0.265	0.723	0.143
51	Information	0.046	0.034	−0.289	−0.098	0.494	−0.086
52	Finance and insurance	0.054	0.059	0.333	0.659	0.634	0.053
53	Rental and leasing	0.010	0.008	0.631	0.486	0.580	−0.000
54	Professional and technical services	0.010	0.016	0.272	0.341	0.728	0.148
56	Administrative and waste services	0.006	0.005	0.506	0.577	0.712	0.132
61	Educational services	0.021	0.026	0.886	0.429	0.793	0.213
62	Health care and social assistance	0.032	0.172	0.205	0.349	0.749	0.169
71	Arts, entertainment, and recreation	0.007	0.011	0.984	0.965	0.664	0.083
72	Accommodation and food services	0.065	0.060	0.498	0.680	0.713	0.133
81	Other services	0.059	0.051	0.382	0.560	0.739	0.159
oo	Housing	0.138	0.150	−0.240	−0.541	0.171	−0.409

Note: Source: CEX, I–O Tables. All moments are aggregated to the 2-digit level and averaged over time (1982–2012). ω_i^{CEX} and ω_i^{PCE} refer to aggregate CEX (resp., PCE) expenditure shares. Expenditure shares correspond to final good expenditures, not value added (thus, agriculture has a tiny expenditure share). For various reasons, CEX and PCE expenditure shares differ; for example, for health care and social assistance, the CEX share is lower because it only captures out-of-pocket expenditures. $\ln(\frac{\omega_i^{Q4}}{\omega_i^{Q1}})$ denotes raw expenditure share differences between high- and low-income households. They closely correspond to whether goods are estimated to be luxuries (i.e., whether their income elasticity γ_i is above one). θ_i^L denotes the labor share and $\theta_i^L - \bar{\theta}^L$ the difference to the economy-wide average. Housing consists of real estate (2002 I–O code 531000) and owner-occupied dwellings (code S00800).

orthogonal to permanent income as proxied by current income, education, and occupation. Here, I discuss these assumptions.

First, there is a concern that the rich spend disproportionately on labor-intensive goods not because of nonhomotheticities, but because of inherent taste differences. To assess that concern, I run an alternative specification that instead controls for education and occupation. The estimated income elasticities γ_i^{ctrl} are plotted against the baseline estimates γ_i in Figure B.4. They closely correlate, with the two outliers being tobacco product manufacturing and book publishers; not surprisingly, the educated smoke less and read more. The correlation between the two sets of estimates is strong: regressing γ_i on γ_i^{ctrl} returns an R^2 of 0.963. The cross-sectional covariance between labor shares and income elasticities falls by 19.6%. However, one does not necessarily want to control for education and occupation when estimating the income elasticities. After all, rising education levels and changing occupational structure are themselves causes and consequences of the growth process, and not orthogonal to it.

Second, there is a concern that the rich pay differential prices, relative to the poor, for the same good. In general, the literature suggests that the rich pay higher prices, as their optimal search effort is lower (Aguiar and Hurst (2007), Kaplan and Menzio (2015)).

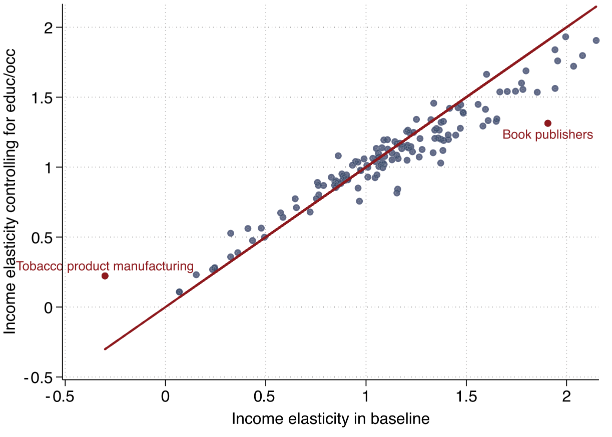

FIGURE B.4.—Income elasticities: Baseline versus controlling for education and occupation.

This suggests that real consumption dispersion is slightly lower than nominal expenditure dispersion; in turn, the variation in income elasticities, and by extension the covariance between income elasticities and labor shares, is slightly underestimated.

I conclude that these potential biases of the income elasticities work in opposite directions, and are unlikely to substantially alter the estimated overall income effect.

B.2. *Production and the Labor Share*

Sensitivity to Individual Sectors. Figure B.5(a) shows that the relation between equipment factor shares and falling labor shares is robust and not driven by any individual sector. The top row visualizes the baseline results (Table I): on average, good-level labor shares drop by 2.5pp; projecting the time trend on average equipment shares accounts for 2.9pp of that decline, implying that the residual time trend increases by 0.4pp according to the point estimate (not significantly different from 0). The other rows demonstrate that no 1-digit sector has an outsized influence on these findings; for example, while outside of manufacturing the overall decline is weaker, the basic finding holds true.

Sensitivity to Timing and Factor Shares versus Capital Intensities. Figure B.5(b) repeats the exercise in Table I, alternatively using the equipment intensity $\bar{\kappa}_i$ instead of the equipment share $\bar{\theta}_i^E$ in the regression in column 3, and/or using initial instead of time-averaged equipment shares/intensities. First, the results show that using equipment intensities $\bar{\kappa}_i$ (effectively controlling for total capital or labor shares) generates the same findings, regardless of whether initial or time-averaged values are used. While the standard errors increase somewhat, the point estimates for the projection on equipment intensity are very similar to the ones for the projection on equipment shares (and by implication, so are the point estimates for the residual time trends). Second, initial equipment factor shares cannot be meaningfully used. This type of regression generates a spurious negative association due to mean reversion in total capital shares (or equivalently, in labor shares), since these are measured as a residual and contain mean-reverting pure profits. High-$\theta_{i,1982}^E$ sectors on average have high initial total capital shares (low initial labor shares); thus, mean reversion biases up the projection of the labor share θ_{it}^L on $t \times \theta_{i,1982}^E$. The time-average $\bar{\theta}_i^E$ does not suffer from this issue, and neither do equipment intensities (since these are de-

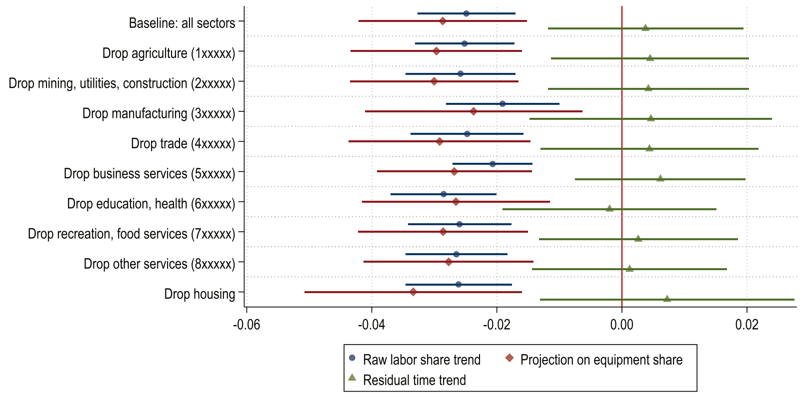

(a) Sensitivity to excluding individual 1-digit sectors

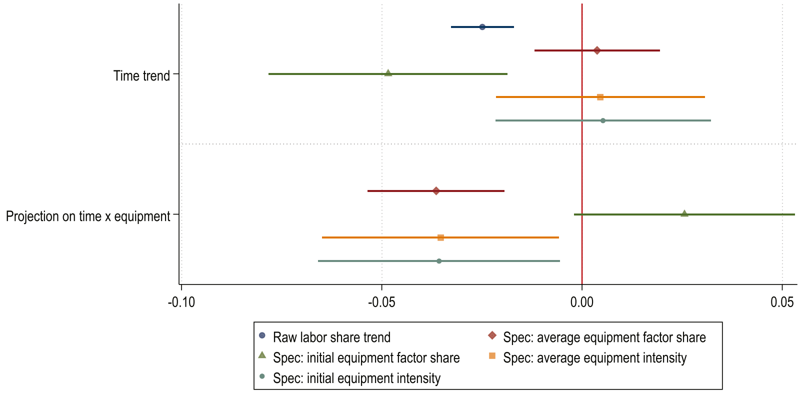

(b) Sensitivity to timing and factor shares versus capital intensities

FIGURE B.5.—Robustness of relation between falling labor share and equipment intensity. Source: BEA I–O Tables (labor shares), NBER-CES Database, BEA FAT (equipment intensities). Both panels correspond to panel-regressions of good-level labor shares with good fixed effects on time trends, as in Table I in the main text. In the upper panel, the baseline repeats Table I, while the remaining specifications drop one 1-digit sector at a time. In the lower panel, the specification with (time-)averaged equipment factor shares ($\bar{\theta}_i^E$) corresponds to Table I, while the other specifications instead interact the time trend with (i) initial equipment factor shares ($\theta_{i,1982}^E$), or (ii) initial equipment intensities ($\kappa_{i,1982}$), or (iii) time-averaged equipment intensities ($\bar{\kappa}_i$).

fined as equipment capital to total capital shares). It is reassuring that all three measures that in theory are not affected by mean reversion generate very similar point estimates.

Good-Factor-Specific Technical Progress. Here, I discuss the potential bias associated with good-factor-specific technological progress. Formally, the residual error term in (16) equals

$$\xi_{it} = (\eta - 1)\big(\bar{\theta}_i^E\big(a_{it}^L - a_{it}^E\big) + \bar{\theta}_i^S\big(a_{it}^L - a_{it}^S\big)\big).$$

Given the estimation strategy with good and time fixed effects and a shift-share regressor, the estimator of η is biased if and only if good-factor-specific technological progress is correlated with equipment shares $\bar{\theta}_i^E$ (likewise, equipment intensities $\bar{\kappa}_i$ in (13)). I focus on

TABLE B.II

SENSITIVITY OF η TO EQUIPMENT-AUGMENTING, GOOD-SPECIFIC TECHNICAL PROGRESS.

	$SD(g_i^E)$						
$Corr(g_i^E, \bar{\theta}_i^E)$	0.001	0.005	0.010	0.020	0.030	0.040	0.050
−1.000	1.396	1.441	1.482	1.039	0.363	0.511	0.647
−0.500	1.391	1.405	1.387	1.277	1.099	1.030	0.947
−0.250	1.387	1.380	1.368	1.285	1.159	1.098	1.042
0.000	1.385	1.373	1.351	1.263	1.204	1.146	1.116
0.250	1.382	1.363	1.328	1.257	1.203	1.164	1.142
0.500	1.380	1.355	1.321	1.264	1.198	1.171	1.146
1.000	1.374	1.335	1.294	1.235	1.194	1.165	1.143

Note: This table reports, for a given value of $SD(g_i^E)$ by column and $Corr(g_i^E, \bar{\theta}_i^E)$ by row, the median IV estimate of η in equation (16) across simulation runs. In each simulation run, g_i^E is drawn from a normal distribution with mean zero, standard deviation $SD(g_i^E)$, and correlation $Corr(g_i^E, \bar{\theta}_i^E)$ with $\bar{\theta}_i^E$. Then the effective change in the log price of equipment capital for good i is computed as $r_{it}^E = r_t^E - (t - 1982)g_i^E$ (instead of $r_{it}^E = r_t^E$ as in the baseline).

the equipment-augmenting term $a_{it}^E = \Delta \ln A_{it}^E$. First, note that if equipment-augmenting progress is stronger in equipment-intensive sectors, as suggested by the theory of directed technical change, then η is biased upwards, but the sign of $(\eta - 1)$ is unbiased. This is because in this case the decline in effective equipment capital costs is understated for equipment-intensive sectors. Thus, the true decline in effective overall capital costs for these sectors was even larger than predicted by measured aggregate capital price variation. In turn, $(\eta - 1)$ is overstated, but of the correct sign (in expectation).

Second, to assess the quantitative importance of this bias, I perform a Monte Carlo simulation. I assume that equipment-augmenting technology is growing at a good-specific rate g_i^E. Table B.II reports the estimates corresponding to equation (16) when taking g_i^E into account. Each cell refers to the median estimate, across simulation runs, for a given value of the annual standard deviation of g_i^E and its correlation with the equipment share $\bar{\theta}_i^E$. For example, $SD(g_i^E) = 0.01$ translates into a standard deviation of 30 log points over 30 years, across goods. For this value and a correlation of 0.5 between g_i^E and $\bar{\theta}_i^E$, the corrected estimate of η is equal to 1.321, while I estimated 1.385 in the baseline. In general, the corrected estimate is decreasing in the dispersion of growth rates. As explained above, as long as this correlation is positive, $(\eta - 1)$ has the correct sign, even if the dispersion of growth rates is very large.[8]

To assess the magnitude of the bias, I use measured variation in TFP growth across 6-digit industries from the NBER-CES Manufacturing Industry Database as a proxy for variation in equipment-augmenting technology. The standard deviation equals 0.431 over 1982–2011; annualizing, $SD(g_i^E) = 0.015$.[9] The correlation with equipment intensity equals 0.169. Using these values, the estimate of η decreases slightly from 1.385 to 1.317.

[8]The fact that the estimates in the middle row corresponding to zero correlation are likewise decreasing in the dispersion of growth rates is a mechanical consequence of effectively adding measurement error in the dependent variable in these simulations, which biases $(\eta - 1)$ toward zero. In this sense, the bias-adjusted estimates in Table B.II should be interpreted as conservative upper bounds on the sensitivity of η to factor-good-specific technical progress.

[9]Across 3-digit industries in the BEA/BLS industry account, the dispersion is slightly lower at 0.013.

APPENDIX C: ADDITIONAL MODEL-BASED FINDINGS

C.1. *Baseline Model Robustness*

Calibration of the Growth Residual. Figure C.1 displays the time series of calibrated labor-augmenting (A_t^L) and factor-neutral (A_t) technology terms in both the benchmark as well as the human capital (HC) calibration. In the benchmark, TFP growth is set to zero ($\Delta A_t = 0$), and all residual growth (i.e., not accounted for by investment-specific technical progress as measured by relative capital prices) is interpreted as labor-augmenting. The HC calibration ties ΔA_t^L to a human capital index, sourced from Penn World Tables 9.0 (Feenstra, Inklaar, and Timmer (2015)), retrieved from FRED (series: HCIYISUSA066NRUG). Because the HC calibration loads less (more) on ΔA_t^L pre-1982 (post-1982), it features a somewhat stronger (weaker) contribution of K–L substitution to the decline in the labor share pre-1982 (post-1982), as shown in Figure 7. Under both calibrations, the growth rate of the labor-augmenting term has decreased considerably around 1980.

Matching the Evolution of Relative Good Prices. In the baseline model, TFP growth is common across goods. Thus, differential price trends across consumption goods are entirely due to differences in factor shares. In the data, there is a lot more price variation. Introducing good-specific TFP growth allows the model to match the price data perfectly. The data source for prices are the BEA's Industry Economic Accounts. In particular, I use annual chained-price indices for gross output by industry on the summary level (71 industries).[10] The findings are hardly affected, as shown in Figure 7 in the main text. The reason is that differential good price trends affect neither the K–L substitution channel, nor the income effect. For intuition, consider first the case of Cobb–Douglas preferences to isolate K–L substitution on the production side. In this case, heterogeneity in TFP growth does not change factor shares at all: Consumers' expenditure shares are fixed. If TFP increases in sector i relative to other sectors, then the consumed quantity of good i is going up, but spending is unchanged. Hence, factor payments are unchanged. Therefore, factor shares are not affected, neither on the sectoral level nor in the aggregate. Moreover, the income effect is not affected by good-specific TFP growth either as long as real income

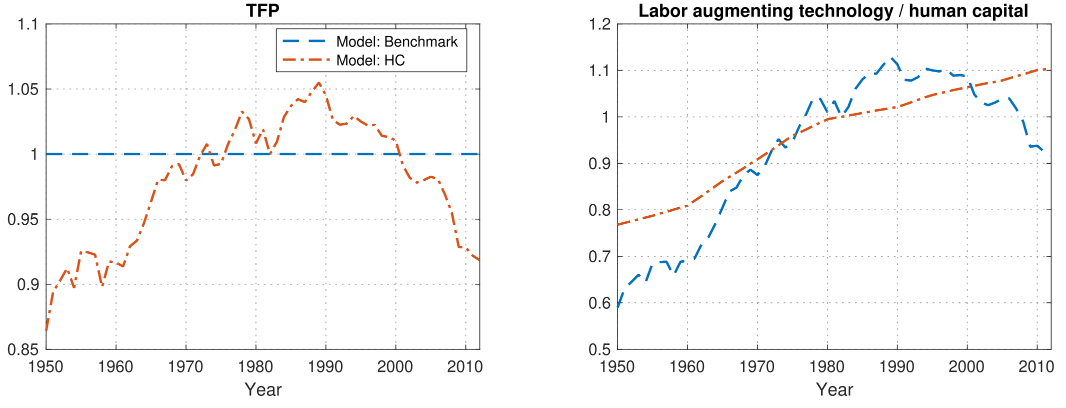

FIGURE C.1.—Calibrated parameters in benchmark and HC calibration.

[10]More detailed data is not available for the time period considered.

growth is unchanged. This leaves merely the contribution of the substitution effect on the consumer side—which is small to begin with—to (potentially) change. Quantitatively, the overall findings are not sensitive to allowing for a richer (and more realistic) pattern of price changes.

C.2. *Investment*

C.2.1. *Investment as a Component of Aggregate Demand*

Table C.I decomposes consumption and investment labor shares into a within-sector component and a "between" or reallocation component. The top right panel reveals that investment spending did not shift toward labor-intensive goods in the data, in contrast to consumption. Note that the reallocation component reflects both substitution as well as income effects. For consumption, income effects are strong and positive. For investment, the notion of an income effect is unclear, as firms are making investment decisions, not households. It is therefore not surprising that the reallocation component is smaller. The within component reveals that investment is concentrated in sectors with faster falling labor shares, relative to consumption.[11]

Calibration of Investment Aggregators. The I–O Tables allow for a breakdown of total private fixed investment into the two types for the sample period 1982–2012, for each final good $i \in I$. The level parameters ω_i^k in equation (19) are chosen such that expenditure shares in the model agree with the data in the base year 1982. Since all prices are normalized to one in the base year (implicitly, the units of goods are chosen appropriately), ω_i^k can be directly equated to the expenditure share of good i in total PFI of type $k = E, S$. I estimate σ^E, similar to the substitution elasticity σ in consumption, from variation in equipment investment shares in response to price changes over time, using price data

TABLE C.I

LABOR SHARE DECOMPOSITION: CONSUMPTION AND INVESTMENT.

	1950–1982			1982–2012		
	Between	Within	Total	Between	Within	Total
Data						
Total (PCE + PFI)				2.1	−9.9	−7.8
Consumption (PCE)				3.0	−8.5	−5.5
Investment (PFI)				0.0	−13.0	−13.1
Model						
Total	3.3	−3.5	−0.2	1.7	−9.2	−7.5
Consumption	4.1	−3.4	0.7	2.9	−8.9	−6.0
Investment	0.1	−4.2	−4.1	0.2	−11.1	−10.9

Note: Source: BEA I–O Tables, own computations. Change in labor share in percentage points, computed on rolling basis (sectoral classifications are time-varying in data). For final demand type $f \in \{\text{Total}, \text{PCE}, \text{PFI}\}$, between component computed as $\sum_{t=2,\dots,T} \sum_{i \in I_t} (\omega_{i,t}^f - \omega_{i,t-1}^f)\theta_{i,t-1}^L$; within as residual: $\sum_{t=2,\dots,T} \sum_{i \in I_t} \omega_{i,t}^f (\theta_{i,t}^L - \theta_{i,t-1}^L)$.

[11]Note that while, relative to consumption, more of investment value added is created in manufacturing, that ratio has been declining as well over time. By now, more than half of investment value added is created in services sectors; see Herrendorf, Rogerson, and Valentinyi (2020) on the implications for modeling structural change.

from the BEA Industry Accounts.[12] The resulting estimate is $\sigma^E = 1.18$ (standard error: 0.04). For structures, I impose $\sigma^S = 1$ (i.e., constant expenditure shares), as the construction sector accounts for more than 80% of structures investment spending throughout.

Model Findings With Investment and Consumption. The bottom panels of Table C.I report results from the model with consumption and investment. Contrary to the baseline model, I am using the BEA's aggregate (PCE) consumption shares here instead of expenditure shares constructed from household micro data (CEX), in order to compare both consumption and investment to comparable aggregate data. Looking at the sample period 1982–2012, the model successfully captures the broad reallocation patterns for both consumption and investment. The model also reflects that investment, relative to consumption, is more heavily concentrated in sectors with faster falling labor shares. Specifically, the final goods that are used for investment purposes are primarily manufacturing and construction goods, which are produced with a relatively high equipment capital intensity. In the model (as in the data), labor shares are falling faster in those sectors because of the steep decline in the price of equipment capital (given that $\eta > 1$).

C.2.2. *Backing Out Nominal Investment Rates in the Baseline Approach*

Here, I describe how the baseline approach implies a time series for nominal investment. To begin with, note that in the baseline approach the constant user cost of capital $\tilde{R}^k$ is a normalization: By virtue of the calibration, the model fits capital shares in the base year $\tau = 1982$ perfectly. A higher $\tilde{R}^k$, for $k = E, S$, decreases the imputed capital-output ratios but does not affect growth rates of factor shares.

The stocks of equipment and structures evolve according to their respective laws of motion (20), where all objects are in efficiency units (quality-adjusted). δ_t^k is the physical depreciation rate, which hardly changes over time, and which I therefore treat as constant.[13] For these implied investment rates, the choice of $\tilde{R}^k$ does matter. A higher $\tilde{R}^k$, by decreasing capital-output ratios, implies lower required investment rates. I choose $\tilde{R}^k$ (more precisely, given δ^k, I choose the sum of the interest rate and expected asset inflation term) such that investment rates in model and data agree over 1950–1982. The subsequent period 1982–2012 can then be used to meaningfully compare model-implied investment rates to the data. Concretely, for any constant $\tilde{R}^k$, the time series of model factor shares yield $(K_t^k)_{t=1950}^{2012}$, since

$$\bar{\theta}_t^k = \frac{\tilde{R}^k P_t^k K_t^k}{Y_t}, \quad k \in \{E, S\}, \tag{C.1}$$

where Y_t is nominal output $\sum_{i \in I} P_{it} Y_{it}$ divided by the consumption deflator (recall that P_t^k is the price of capital relative to consumption). In turn, the law of motion (20) implies real investment I_t^k, which in turn can be translated into nominal investment rates i_t^k:

$$i_t^k = \frac{P_t^k I_t^k}{Y_t}, \quad k \in \{E, S\}. \tag{C.2}$$

[12]The BEA price data is not fully quality-adjusted, requiring to add an extra TFP term A_t^k. Without this extra term, the model price of the equipment aggregate, relative to the consumption aggregate, would not decline as fast as in the targeted data series (DiCecio (2009)).

[13]See Cummins and Violante (2002) for a discussion of economic versus physical depreciation rates, and corroborating evidence for constancy of δ_t^E. Based on BEA Fixed Assets Table depreciation data and removing obsolescence due to the change in the relative price of the asset, I find $\delta^E = 0.098$ and $\delta^S = 0.027$.

To sum up, the baseline approach followed in this paper implies time series of nominal investment rates i_t^k through the equations (C.1), (20), (C.2); one for equipment and one for structures. Observe that the i_t^k are, for each element, monotonically decreasing in $\tilde{R}^k$. Thus, for $k = E, S$, there is a unique $\tilde{R}^{k,\star}$ such that i_t^k, averaged over 1950–1982, matches the data equivalent. Given this choice of $\tilde{R}^{k,\star}$, model investment rates can be meaningfully compared to the data over 1982–2012.

C.2.3. *Calibration of I-Model*

I impose that the I-Model replicates the baseline model equilibrium in the initial year 1950 (in particular, $\tilde{R}^k_{1950} = \tilde{R}^{k,\star}$). While other model parameters are borrowed from the baseline model, I recalibrate the time series of labor-augmenting technology A_t^L such that real consumption growth still matches the data. By construction, the I-Model agrees with the baseline on average over 1950–1982, when it predicts a stable aggregate labor share (because over that time period, investment rates in the baseline agree with the data).

C.2.4. *Increasing Capital Share and Increasing $\tilde{R}^k_t$ in the I-Model*

Why does the I-Model feature both an increasing capital share as well as an increasing user cost per dollar of capital $\tilde{R}^k_t$? These points are easier to understand through the more familiar equations of a one-sector model with just one type of capital K. Let $k_t \equiv \frac{K_t}{Y_t}$ denote the physical capital-output ratio, and let P_t^K denote the relative price of the aggregate investment good. Then the FOC for capital demand is

$$\tilde{R}_t P_t^K = \alpha^{\frac{1}{\eta}} k_t^{-\frac{1}{\eta}}$$

and the equilibrium capital share is

$$\theta_t^K \equiv \frac{\tilde{R}_t P_t^K K_t}{Y_t} = \alpha^{\frac{1}{\eta}} k_t^{\frac{\eta-1}{\eta}} = \alpha\big(\tilde{R}_t P_t^K\big)^{1-\eta}. \tag{C.3}$$

The law of motion of capital (20) can be rewritten in terms of k_t and the growth rate of physical capital, $g_{t+1}^K \equiv \frac{K_{t+1}}{K_t} - 1$, as

$$\big(g_{t+1}^K + \delta\big) k_t = \frac{i_t}{P_t^K}. \tag{C.4}$$

As P_t^K is decreasing at a faster rate post-1982, from (C.4) it is apparent that the growth rate of physical capital increases given a roughly constant nominal investment rate i_t. This increases k_t, and since $\eta > 1$, from (C.3) it follows that the capital share increases. The simple one-sector model does of course not feature the counteracting forces of nonhomothetic demand; still, the intuition applies.

Second, as the capital share increase is more modest in the I-Model and $\tilde{R}^k$ is constant in the baseline, (C.3) illustrates why $\tilde{R}^k_t$ has to increase over time in the I-Model for capital market clearing.

C.3. *Increasing Income and Consumption Inequality*

The model abstracts from changes in consumer heterogeneity, focusing instead on mean income and consumption growth. Under nonhomothetic demand, the distribution

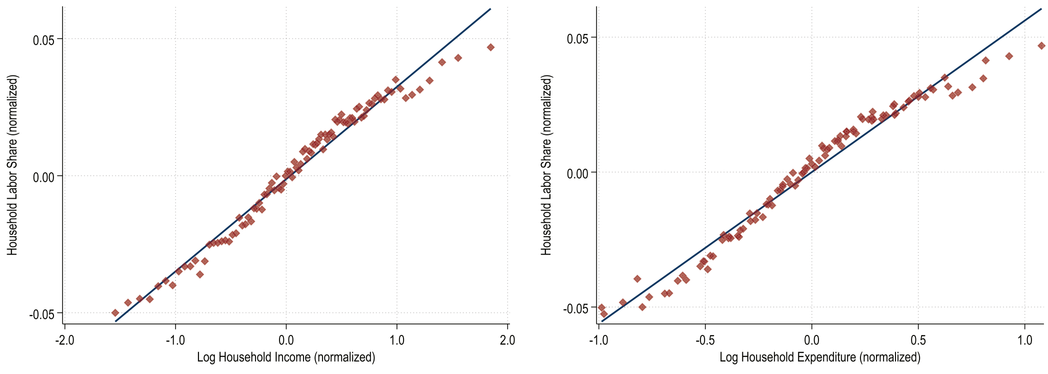

(a) Income (b) Consumption

FIGURE C.2.—Log-linearity of household consumption labor shares. Source: CEX, I–O Tables. Households grouped in percentile-year bins according to after-tax household income. Consumption, income, and the labor share embodied in households' consumption baskets are first demeaned within each year; then averages over all years (1980–2015) are reported. For the income graph on the left panel, the bottom five percentiles are truncated, as they have very low income (all percentiles are reported for consumption).

of total consumption expenditures across households generally matters for the composition of aggregate demand. Moreover, consumption inequality has increased over the past few decades. However, the fact that households' consumption labor shares are monotonically increasing in income does not necessarily imply that more consumption inequality increases the aggregate labor share. Instead, whether this implication is true depends on the relative extent of nonhomotheticicities in different parts of the consumption distribution.[14]

Figure C.2 plots households' consumption labor shares against log income and log total expenditure. The shape of these relations is close to linear, except for the top percentiles. Under exact linearity, it is possible to derive an analytic expression for the partial equilibrium change in the aggregate labor share in response to changes in a log-normal consumption distribution: Assume that

$$\theta_h^L = c_0 + c_1 \ln(E_h),$$

where θ_h^L is the labor share of household h, E_h its total expenditure, and c_0, c_1 are scalars. Let $\ln(E_h) \sim N(\mu - \frac{\sigma^2}{2}, \sigma^2)$, so that $\mathbb{E}[E_h] = \exp(\mu)$. Then

$$\bar{\theta}^L \equiv \frac{\mathbb{E}[E_h \theta_h^L]}{\mathbb{E}[E_h]} = \frac{\mathbb{E}[E_h(c_0 + c_1 \ln(E_h))]}{\mathbb{E}[E_h]} = c_0 + c_1 \frac{\mathbb{E}[E_h \ln(E_h)]}{\mathbb{E}[E_h]} = c_0 + c_1\left(\mu + \frac{\sigma^2}{2}\right).$$

[14]To see this, note that if demand were derived from Stone–Geary preferences (and all consumers were above the subsistence level), then mean-preserving changes in the income distribution would have no effect at all on aggregate demand. In this special case, the two opposing effects cancel out exactly: on the one hand, the rich become even richer, which pushes up the aggregate labor share; on the other hand, the poor become poorer, which makes them spend dis-proportionately less on labor-intensive goods, relative to the decrease in their total consumption expenditure.

The last equality follows since $\mathbb{E}[E_h] = \exp(\mu)$, and using the notation $y_h \equiv \ln(E_h)$,

$$\begin{aligned}
\mathbb{E}\big[E_h \ln(E_h)\big] &= \mathbb{E}\big[\exp(y_h) y_h\big] \\
&= \int \exp(y) y \frac{1}{\sqrt{2\sigma^2\pi}} \exp\left(-\frac{\big(y-\mu+\sigma^2/2\big)^2}{2\sigma^2}\right) dy \\
&= \int y \frac{1}{\sqrt{2\sigma^2\pi}} \exp\left(y - \frac{y^2 - 2\mu y + y\sigma^2 - \mu\sigma^2 + \mu^2 + \sigma^4/4}{2\sigma^2}\right) dy \\
&= \int y \frac{1}{\sqrt{2\sigma^2\pi}} \exp\left(-\frac{y^2 - 2y\big(\mu+\sigma^2/2\big) + \big(\mu+\sigma^2/2\big)^2 - 2\mu\sigma^2}{2\sigma^2}\right) dy \\
&= \exp(\mu) \int y \frac{1}{\sqrt{2\sigma^2\pi}} \exp\left(-\frac{\big(y-\big(\mu+\sigma^2/2\big)\big)^2}{2\sigma^2}\right) dy \\
&= \exp(\mu)\big(\mu+\sigma^2/2\big).
\end{aligned}$$

Real personal consumption expenditures grew by 87% or 63 log points over the period 1982–2007, per capita. Over the same time period, estimates for the increase in the variance of log consumption range from 6 to 18 points.[15] Using these numbers, factoring in rising inequality increases the strength of the positive income effects on the labor share by $\frac{\Delta\sigma^2/2}{\Delta\mu} = 5 - 14\%$. Since linearity breaks down at the very top according to Figure C.2, I view these numbers as upper bounds. I conclude that while this channel would strengthen the positive income effects on the labor share, it is an order of magnitude smaller than the effect of rising mean income. In light of this, and factoring in the uncertainty concerning the extent of increasing inequality and the exact shape of the household expenditure-labor share relation, the analysis in this paper abstracts from changes in consumer heterogeneity over time.

APPENDIX D: PROOF OF PROPOSITION 1

PROOF: Drop time indices for clarity. To simplify notation, define $g_L \equiv d\ln A^L\bar{L}$ and $g_K \equiv d\ln A^K\bar{K}$ as the respective log changes of labor and capital in efficiency units. Also define $g_W \equiv d\ln\frac{W}{A^L}$ and $g_R \equiv d\ln\frac{R}{A^K}$ as the log changes in their respective factor prices, again in efficiency units. Let $g_A = d\ln A$. Let the wage rate, in efficiency units, be the numeraire: $g_W = 0$. As in the main text, $g \equiv d\ln\frac{E}{P}$ denotes the overall real growth rate.

The proof uses the market clearing condition for labor to solve for g_R in terms of the fundamentals g_L and g_K. In turn, the equilibrium change in the aggregate labor share can be expressed in terms of g_L, g_K and g (where the latter is a function of g_L, g_K, g_A).

First, the market clearing condition for labor can be written as

$$W\bar{L} = E\left(\sum_{i\in I}\omega_i\theta_i^L\right), \tag{D.1}$$

[15]See Attanasio and Pistaferri (2016) for an overview of the literature on inferring consumption inequality. The consensus has shifted toward the higher end of estimates; see, in particular, Aguiar and Bils (2015).

labor income equals expenditure on labor. Take logs and totally differentiate (D.1):

$$d\ln(W\bar{L}) = \underbrace{d\ln\left(\frac{W}{A^L}\right)}_{\equiv g_W=0} + \underbrace{d\ln\left(A^L\bar{L}\right)}_{\equiv g_L} = d\ln E + d\ln\left(\sum_{i\in I}\omega_i\theta_i^L\right)$$

$$\Rightarrow \quad g_L = d\ln E + \frac{\sum_{i\in I}\omega_i\theta_i^L\left(d\ln\omega_i + d\ln\theta_i^L\right)}{\bar{\theta}^L}. \tag{D.2}$$

Observe that

$$d\ln E = d\ln(W\bar{L} + R\bar{K}) = \bar{\theta}^L g_L + \left(1-\bar{\theta}^L\right)(g_R + g_K). \tag{D.3}$$

Furthermore, $d\ln P_i = (1-\theta_i^L)g_R - g_A$ by Shephard's lemma. Hence,

$$d\ln P \equiv \sum_{i\in I}\omega_i\, d\ln P_i = \left(1-\bar{\theta}^L\right)g_R - g_A. \tag{D.4}$$

Using (D.3) and (D.4), the overall growth rate is as in the one-sector growth model:

$$g \equiv d\ln\frac{E}{P} = g_A + \bar{\theta}^L g_L + \left(1-\bar{\theta}^L\right)g_K.$$

Substituting in for $d\ln\frac{P_i}{P}$ and $d\ln\frac{E}{P}$ in (7),

$$d\ln\omega_i = (1-\sigma)\left(\bar{\theta}^L - \theta_i^L\right)g_R + (\gamma_i - 1)g. \tag{D.5}$$

Note that (4), in the special case of a single type of capital, and given $g_W = 0$, implies

$$d\ln\theta_i^L = (\eta-1)\left(1-\theta_i^L\right)g_R. \tag{D.6}$$

Substituting for $d\ln E$ from (D.3), $d\ln\omega_i$ from (D.5), and $d\ln\theta_i^L$ from (D.6) in (D.2):

$$g_L = \bar{\theta}^L g_L + \left(1-\bar{\theta}^L\right)(g_R + g_K)$$
$$+ \frac{\sum_{i\in I}\omega_i\theta_i^L\left((1-\sigma)\left(\bar{\theta}^L - \theta_i^L\right)g_R + (\gamma_i-1)g + (\eta-1)\left(1-\theta_i^L\right)g_R\right)}{\bar{\theta}^L}.$$

Rearranging this equation, we can write it as

$$g_L - g_K = \tilde{\eta}g_R + g\frac{\mathrm{Cov}\left(\theta_i^L, \gamma_i\right)}{\bar{\theta}^L\left(1-\bar{\theta}^L\right)}, \tag{D.7}$$

where $\mathrm{Cov}(\theta_i^L, \gamma_i) = \sum_{i\in I}\omega_i(\gamma_i\theta_i^L - \bar{\theta}^L)$ and

$$\tilde{\eta} = \frac{\sigma\mathbb{V}\left[\theta_i^L\right] + \eta\mathbb{E}\left[\theta_i^L\left(1-\theta_i^L\right)\right]}{\bar{\theta}^L\left(1-\bar{\theta}^L\right)}.$$

Note that $\tilde{\eta}$ is indeed a convex combination of σ and η, since $\mathbb{V}[\theta_i^L] + \mathbb{E}[\theta_i^L(1-\theta_i^L)] = \mathbb{E}[(\theta_i^L)^2] - (\bar{\theta}^L)^2 + \bar{\theta}^L - \mathbb{E}[(\theta_i^L)^2] = \bar{\theta}^L(1-\bar{\theta}^L)$. If all sectors use capital and labor in the same proportion, then $\tilde{\eta} = \eta$. On the other hand, if $\theta_i^L \in \{0, 1\}\ \forall i \in I$ (the polar opposite case of maximal variability in labor shares across sectors), then $\tilde{\eta} = \sigma$.

Plugging in for $d \ln E$ from (D.3), the change in the aggregate labor share is given by

$$d \ln \bar{\theta}^L = d \ln \frac{W\bar{L}}{E} = g_L - d \ln E = (1 - \bar{\theta}^L)(g_L - g_K - g_R). \tag{D.8}$$

Finally, substituting for g_R from (D.7) in (D.8) yields

$$\begin{aligned} d\bar{\theta}^L &= \bar{\theta}^L\, d \ln \bar{\theta}^L = \bar{\theta}^L(1-\bar{\theta}^L)\left(g_L - g_K - \frac{1}{\tilde{\eta}}\left(g_L - g_k - g\frac{\mathrm{Cov}(\theta_i^L, \gamma_i)}{\bar{\theta}^L(1-\bar{\theta}^L)}\right)\right) \\ &= \frac{\tilde{\eta}-1}{\tilde{\eta}}\bar{\theta}^L(1-\bar{\theta}^L)(g_L - g_K) + \frac{g}{\tilde{\eta}}\mathrm{Cov}(\theta_i^L, \gamma_i) \\ &= \frac{\tilde{\eta}-1}{\tilde{\eta}}\bar{\theta}^L(1-\bar{\theta}^L)\left(d \ln A^L\bar{L} - d \ln A^K\bar{K}\right) + \frac{g}{\tilde{\eta}}\mathrm{Cov}(\theta_i^L, \gamma_i), \end{aligned}$$

which proves (8). *Q.E.D.*

REFERENCES

AGUIAR, MARK, AND MARK BILS (2015): "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review*, 105, 2725–2756. [5,20]

AGUIAR, MARK, AND ERIK HURST (2007): "Life-Cycle Prices and Production," *American Economic Review*, 97, 1533–1559. [11]

ATTANASIO, ORAZIO P., AND LUIGI PISTAFERRI (2016): "Consumption Inequality," *Journal of Economic Perspectives*, 30, 3–28. [20]

CABALLERO, RICARDO J., EMMANUEL FARHI, AND PIERRE-OLIVIER GOURINCHAS (2017): "Rents, Technical Change, and Risk Premia Accounting for Secular Trends in Interest Rates, Returns on Capital, Earning Yields, and Factor Shares," *American Economic Review*, 107, 614–620. [6]

CUMMINS, JASON G., AND GIOVANNI L. VIOLANTE (2002): "Investment-Specific Technical Change in the US (1947–2000): Measurement and Macroeconomic Consequences," *Review of Economic Dynamics*, 5, 243–284. [17]

DE LOECKER, JAN, EECKHOUT JAN, AND GABRIEL UNGER (2020): "The Rise of Market Power and the Macroeconomic Implications," *The Quarterly Journal of Economics*, 135, 561–644. [7]

DICECIO, RICCARDO (2009): "Sticky Wages and Sectoral Labor Comovement," *Journal of Economic Dynamics and Control*, 33, 538–553. [17]

FEENSTRA, ROBERT C., ROBERT INKLAAR, AND MARCEL P. TIMMER (2015): "The Next Generation of the Penn World Table," *American Economic Review*, 105, 3150–3182. [15]

GARNER, THESIA I., GEORGE JANINI, WILLIAM PASSERO, LAURA PASZKIEWICZ, AND MARK VENDEMIA (2006): "The CE and the PCE: A Comparison," *Monthly Labor Review*, 129, 20–46. [5]

GOLLIN, DOUGLAS (2002): "Getting Income Shares Right," *Journal of Political Economy*, 110, 458–474. [1]

HERRENDORF, BERTHOLD, RICHARD ROGERSON, AND ÁKOS VALENTINYI (2020): "Structural Change in Investment and Consumption—A Unified Analysis," *The Review of Economic Studies*, 88, 1311–1346. [16]

HOROWITZ, KAREN J., AND MARK A. PLANTING (2014): "Concepts and Methods of the U.S. Input–Output Accounts," U.S. Department of Commerce, Bureau of Economic Analysis. [3]

HUBMER, JOACHIM, AND PASCUAL RESTREPO (2021): "Not a Typical Firm: The Joint Dynamics of Firms, Labor Shares, and Capital–Labor Substitution," Working Paper 28579, National Bureau of Economic Research. [7]

KAPLAN, GREG, AND GUIDO MENZIO (2015): "The Morphology of Price Dispersion," *International Economic Review*, 56, 1165–1206. [11]

LEVINSON, ARIK, AND JAMES O'BRIEN (2019): "Environmental Engel Curves: Indirect Emissions of Common Air Pollutants," *Review of Economics and Statistics*, 101, 121–133. [4]
VALENTINYI, AKOS, AND BERTHOLD HERRENDORF (2008): "Measuring Factor Income Shares at the Sector Level," *Review of Economic Dynamics*, 11, 820–835. [1]

Co-editor Charles I. Jones handled this manuscript.

Manuscript received 11 June, 2020; final version accepted 14 October, 2022; available online 26 October, 2022.