The Race Between Preferences and Technology

Joachim Hubmer *
Yale University

First version: November 6, 2018
This version: December 22, 2018

Job Market Paper

Link to most current version:
`https://sites.google.com/site/joachimhubmer/jmp`

Abstract

This paper argues that a unified analysis of consumption and production is required to understand the long-run behavior of the labor share of income in the United States. First, using household data on the universe of consumer spending, I document that higher-income households spend relatively more on labor-intensive goods and services as a share of their total consumption. Interpreted as the result of non-homothetic preferences, this fact implies that economic growth increases the labor share through an income effect. Second, using disaggregated good-level data on factor shares and capital intensities, I estimate that capital and labor are gross substitutes. Consequently, investment-specific technical change, manifesting itself in the form of a well-documented decline in the relative price of equipment capital, reduces the labor share. Given the estimated elasticities, I show that a parsimonious neoclassical model quantitatively matches the observed low-frequency movement in the aggregate labor share since the 1950s, both its relative stability until about 1980 and its decline thereafter. Until the early 1980s, the income effect, working through non-homothetic preferences, offset capital-labor substitution. Subsequently, accelerating investment-specific technical change, leading to increasing substitution of capital for labor, began to dominate the income effect.

*I am indebted to my advisors Giuseppe Moscarini, Tony Smith, and Per Krusell for their continuous guidance and support throughout this project. I am grateful to Costas Arkolakis, Bill Brainard, Zhen Huo, Ilse Lindenlaub, Michael Peters, Pascual Restrepo, Aleh Tsyvinski, Fabrizio Zillibotti, as well as to seminar participants at the Yale Macro Lunch, and at the Minneapolis Fed Junior Scholar Conference, for their comments and feedback. I also thank James O'Brien for generously sharing a concordance between CEX categories and Input-Output Tables with me. All errors are my own. *Email*: joachim.hubmer@yale.edu.

1 Introduction

We are witnessing an era of rapid technological advances, manifesting itself in new, better, or cheaper machines (stark examples are robots and artificial intelligence). As these new technologies that are embodied in capital goods diffuse in the economy, fears about the future of labor opportunities abound. Indeed, the share of labor in national income has been declining in the U.S. as well as globally over the past few decades (Elsby, Hobijn and Sahin, 2013, Karabarbounis and Neiman, 2014). These fears have re-emerged continually since the beginning of the Industrial Revolution. Yet, labor has not become redundant. Kaldor (1961) famously mentioned the observed stability of the labor share as one of the stylized facts of economic growth.

This paper studies the evolution of the U.S. labor share in the post-war era (Panel (a) of Figure 1).[1] Throughout the entire period, we witnessed remarkable technological progress in the sectors that produce equipment capital. Following the literature, I use the quality-adjusted real equipment price (Panel (b) of Figure 1) as a measure of investment-specific technical change.[2] Note that this price decline accelerated substantially in the early 1980s. As machines become relatively cheaper, they are substituted for labor. However, two obstacles seemingly derail the narrative that labor shares are falling because of technological changes embodied in capital goods.

First, for labor shares to decline in response to falling capital costs, capital and labor have to be gross substitutes in production, contrary to the prevailing consensus in the literature.[3] My first contribution is to show that, when adopting a disaggregated good-level perspective of the U.S. economy and accounting for input-output linkages, capital and labor are indeed gross substitutes. Key for this empirical result is my first new stylized fact: in the cross-section of goods, labor shares are falling in proportion to the equipment-intensity of capital (Panel (c) of Figure 1).

Second, this narrative in isolation has difficulty accounting for the stability of the aggregate labor share prior to the 1980s. My second contribution is to identify a countervailing force. The very changes in technology that improve machines are also a source of economic growth. Rising levels of real income lead to different spending patterns. Crucially, I document that richer households spend relatively more on labor-intensive goods and services as a fraction of total expenditure (Panel (d) of Figure 1). Thus, holding constant relative prices, economic growth increases the aggregate labor share through an income effect.

My third contribution is to show that an otherwise parsimonious neoclassical model with non-homothetic preferences quantitatively matches the observed low-frequency movement in the

[1]Because there is some ambiguity concerning the treatment of proprietors' income, I report both the BLS' headline labor share measure and the narrower payroll share, and note that the totality of the evidence suggests that the labor share was roughly stable until the early 1980s and declined subsequently, possibly accelerating later on. For all time series in Figure 1, I added separate linear time trends for the pre-1982 and the post-1982 time period for illustrative purposes only. In the model analysis, I relate to the raw annual time series.

[2]The series is constructed by DiCecio (2009), building on earlier work by Gordon (1990) and Cummins and Violante (2002). Hulten (1992) and Greenwood, Hercowitz and Krusell (1997) are seminal references for investment-specific technical change.

[3]Karabarbounis and Neiman (2014) is an important exception to this consensus. Across countries, they similarly find that falling investment good prices correlate with falling labor shares, and conclude that investment-specific technical change explains half of the global decline in the labor share.

Figure 1: Two aggregate trends and two facts about the cross-section

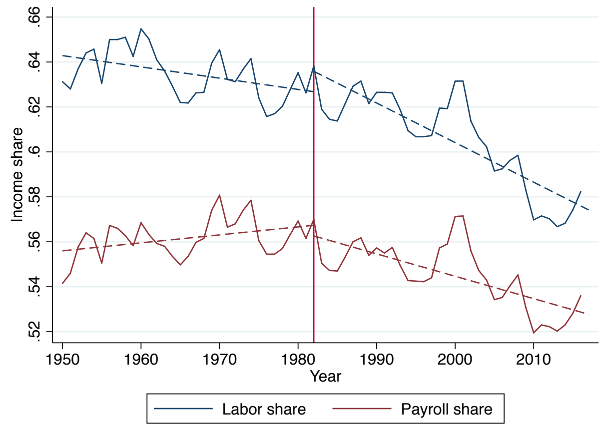

(a) Aggregate U.S. labor share

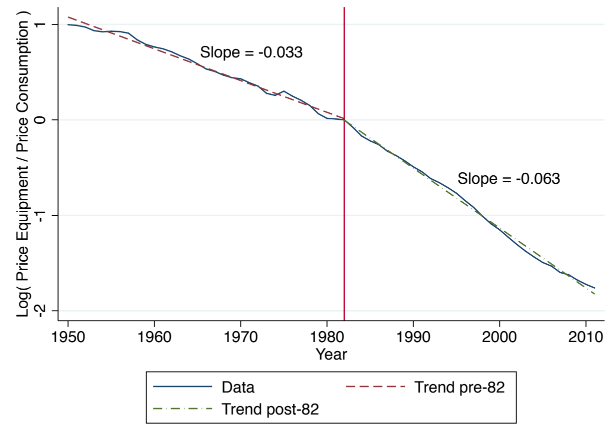

(b) Relative price of equipment and software

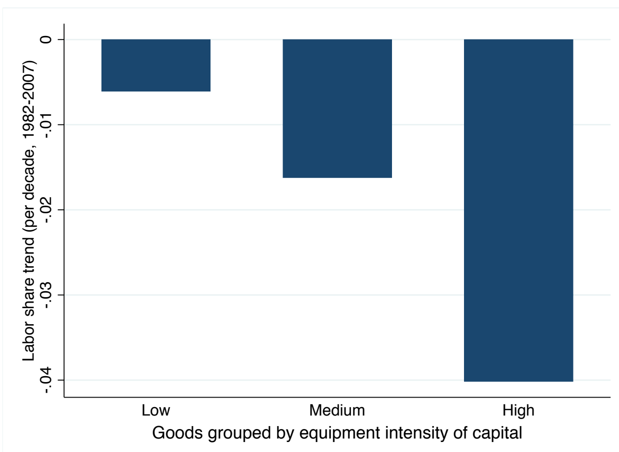

(c) Trends in good-level labor shares

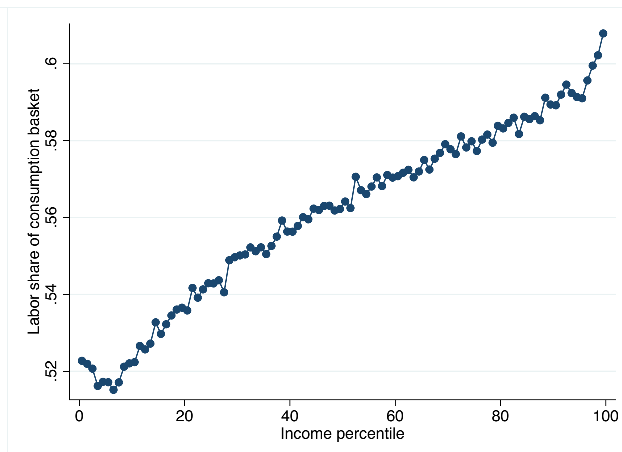

(d) Household labor shares

Data sources: (a) BLS; (b) DiCecio (2009); (c) BEA I-O Tables, BEA FAT, NBER-CES Manufacturing Database, own computations; (d) CEX, BEA I-O Tables, own computations.

aggregate labor share, both its relative stability until about 1980 *and* its decline thereafter.

I first examine the question theoretically in a neoclassical general equilibrium framework. I show that the response of the aggregate labor share to different forms of economic growth can be decomposed into two additive components: The first is a substitution effect, operating both on the production side via direct capital-labor substitution, as well as indirectly via a reallocation of consumption in response to changing prices. This substitution effect is proportional to the bias of growth towards capital, and depends on elasticities of substitution in production as well as in consumer demand. The second component represents an income effect, and is proportional to the overall rate of economic growth multiplied by the cross-sectional covariance between sectoral labor shares and income elasticities. If this covariance is positive and the relevant aggregate substitution elasticity above one, then the aggregate labor share is stable if growth exhibits a moderate capital bias, while it declines if the capital bias is strong.

Guided by the theoretical analysis, I estimate the key elasticities. First, using the Bureau of Economic Analysis' (BEA) Detailed Input-Output (I-O) Tables, I construct a panel dataset of labor shares at the final good level, reflecting all upstream value added.[4] The I-O Tables provide a complete, fine-grained, picture of the U.S. economy, covering around four hundred industries and as many goods. Linking this panel to the U.S. Consumer Expenditure Survey (CEX), which contains household data on the universe of consumer spending, I document that richer households spend more on labor-intensive goods as a fraction of total expenditure. This pattern holds over the entire sample period (1980–2015), implying that the covariance between income elasticities and labor shares is positive throughout. Interpreted in a framework with stable preferences, this non-homotheticity implies that any form of economic growth increases the aggregate labor share through an income effect.

Next, I turn to estimating the capital-labor elasticity of substitution in production, again relying on the panel of good-level labor (and capital) shares derived from the BEA's I-O Tables. The estimation strategy is based on the assumption that the observed secular decline in the quality-adjusted price of equipment and software capital reflects exogenous technical progress. I document that the fall in labor shares, across goods, has been proportional to the equipment intensity of capital. To construct the latter variable, I use data from the BEA's Fixed Asset Tables and the NBER-CES Manufacturing Industry Database, as well as data on public firms from Compustat. Because in the cross-section of goods, falling labor shares correlate with falling capital costs, the estimated capital-labor elasticity is significantly above one. In other words, capital and labor are gross substitutes, implying that declining capital prices trigger falling labor shares, as firms optimally shift expenditure towards the cheaper production factor. The estimated elasticity represents the shift from labor to capital in response to a fall in the rental rate of capital, relative to the wage rate, along the full value chain, including all upstream production.

Armed with these two sets of estimates, I analyze the evolution of the U.S. labor share since the 1950s through the lens of the model. For this exercise, I assume that the key consumer demand and technology elasticities have been stable over time. The model quantitatively matches the observed low-frequency movement in the aggregate labor share, both its relative stability until about 1980 and its decline thereafter. Up to the early 1980s, the substitution of capital for labor in production was moderate, and, as it turns out, entirely offset by the positive income effect. Later on, as investment-specific technical change accelerated, capital-labor substitution became the dominating force.[5] Factoring in the non-homotheticities in demand is crucial, as homothetic versions of the model fail to rationalize the data.

An emergent strand of the literature argues that increasing market power has led to rising markups and profits (Barkai, 2017).[6] De Loecker and Eeckhout (2017) use Compustat data to

[4] Throughout this paper, I will use goods as a shorthand for goods and services.

[5] This acceleration of the fall in real equipment prices has also been noted by Fisher (2006).

[6] Barkai (2017) constructs the capital share as the product of the capital stock times the required rate of return, and defines the profit share as the residual of labor and capital shares. This contrasts with the baseline approach in

estimate a strong increase in markups over marginal costs since the 1980s, driven by an increase in the ratio of variable profits to sales. Hall (2018) documents a (smaller) increase on the industry-level using KLEMS data. If labor shares of equipment-intensive sectors declined not because of capital-labor substitution but because of an increase in economic profits, then my estimate of the capital-labor elasticity would be upward biased. However, both for Hall's as well as for Compustat-based markup estimates, I find that rising markups only weakly predict falling labor shares in the cross-section, and that they are largely orthogonal to equipment-intensity. Therefore, and importantly, my estimate of the capital-labor elasticity is virtually unaffected when controlling for time-sector specific markups. I interpret this finding as support for the capital-labor substitution channel instead of an increase in economic profits.[7] Further assessing robustness, I investigate the potential of good-factor-specific technical progress and good-specific substitution elasticities to bias my estimate of the capital-labor elasticity. While I find that these may affect the precise value of the estimate, I conclude that it is unlikely that the true elasticity is smaller than one.

The baseline model abstracts from several layers of complexity, which I did not find to be crucial and which I therefore address separately. First, it does not feature increasing household heterogeneity. In general, the distribution of consumer spending across households matters if demand is non-homothetic. However, I find that while factoring in rising dispersion would slightly increase the estimated aggregate income effect, its contribution is at least an order of magnitude smaller than the one of rising mean income, which this paper focuses on.

Second, in the baseline model I assume that all value added is created domestically. While the import ratio of the U.S. economy has increased substantially, I find no evidence for this channel to affect the results in this paper. On the one hand, since import shares are relatively similar across household income groups, this simplifying assumption does not affect the estimated income effect. On the other hand, I find that the estimate of the capital-labor elasticity is not confounded by differential import exposure. Moreover, the factor content of the basket of imported goods is in fact very similar, both in terms of levels as well as in terms of changes, to the one of exports.

Finally, I address the extent to which the findings are consistent with aggregate investment data. In the baseline exercise, I treat factor prices of capital as exogenous inputs to the model, and back out the resulting capital stocks. This procedure implies a time series of nominal investment rates that can be compared to the data. While in the data the ratio of private investment to the sum of private investment and consumption has fluctuated without any apparent trend around its long-run mean of 21%, the baseline model implies an increase to 26%. This discrepancy motivates studying an alternative model version that treats the investment data as the truth. In that alternative

this paper, which defines the capital share as the residual of the labor share, implicitly assuming that either part of the capital stock is unmeasured and / or the required rate of return is higher. See Karabarbounis and Neiman (2018) for a discussion of the implications of alternative strategies of how to deal with what they label "factorless income".

[7] To rationalize this finding, I note that the evidence on rising market power is mixed. Rossi-Hansberg, Sarte and Trachter (2018) document that while product-market concentration has increased on the national level, the *opposite* is true on the local level. Moreover, the aggregate markup derived from Compustat firms has been shown to be quite sensitive to the choice of aggregation weights (Edmond, Midrigan and Xu, 2018), to the definition of variable costs, and to assumptions on the representativeness of public firms (Traina, 2018).

model, the fall in the labor share is less dramatic, though still quite large at 3.7%, compared to 5.3% in the baseline. It is perhaps surprising that the capital share increases even if the nominal investment rate (and capital-output ratio) are relatively stable over time. However, even with flat nominal investment, investment-specific technical change leads to capital deepening in real terms. Combined with an above-unitary capital-labor elasticity, the capital share increases.

The paper proceeds as follows. Section 2 relates to the literature. Section 3 contains the modeling framework and analytically characterizes the forces that shape the evolution of the labor share in a neoclassical model with non-homothetic preferences. Section 4 introduces consumption micro data as well as disaggregated production data, provides descriptive evidence, and discusses identification and estimation. Section 5 presents the main results. Section 6 discusses details and context on estimating the capital-labor elasticity, and Section 7 other robustness and extensions. Finally, Section 8 concludes. The appendix contains all proofs, as well as further empirical and quantitative results.

2 Related literature

First and foremost, this paper contributes to the voluminous literature on the evolution of the labor share in general, and its decline in the past few decades in particular. Elsby et al. (2013) document the decline for the U.S., while Karabarbounis and Neiman (2014) show that it is a global phenomenon. The exact magnitude of the decline is still debated due to measurement issues such as the treatment of the labor portion of proprietor's income (Gollin, 2002, Elsby et al., 2013), intangible capital (Koh, Santaeullia-Llopis and Zheng, 2016), and housing (Rognlie, 2015). There is, however, a consensus that the labor share has indeed been falling. Many explanations have been put forward, including: increased openness to international trade (Elsby et al., 2013); an increase in concentration that causes increasing profit rates (Barkai, 2017, Autor, Dorn, Katz, Patterson and Reenen, 2017) and relatedly increasing markups (De Loecker and Eeckhout, 2017); automation (Acemoglu and Restrepo, 2018, 2017), as well as more generally capital-labor substitution that is triggered either by declining investment good prices (Karabarbounis and Neiman, 2014, Eden and Gaggl, 2018) or by capital accumulation itself (Piketty, 2014).

Relative to the majority of studies in that literature, the focus on accounting for the behavior of the labor share over the entire post-war period in the U.S. is novel. Karabarbounis and Neiman (2018) is a closely related paper insofar as they also go beyond the period of declining labor shares (post-1980) to argue that some of the proposed explanations are seemingly at odds with earlier data. Conceptually, I contribute to the literature by highlighting the need to study production and consumption in a joint framework to understand the long run—in particular, to rationalize not only the recent decline, but also why the labor share used to be stable. Empirically, I contribute by estimating the key elasticities in consumption and production using an array of disaggregated production data and household consumption data. Quantitatively, I contribute by showing that given these estimated elasticities, the neoclassical model matches the observed low-frequency movement in the aggregate U.S. labor share in the post-war era if (and only if) it is extended to many sectors and non-homothetic demand.

Second, this paper also contributes to the literature on estimating the capital-labor elasticity. In the interest of space, I refer to summaries of the previous literature in Chirinko (2008) and León-Ledesma, McAdam and Willman (2010). Conceptually, I build on Karabarbounis and Neiman (2014), who also focus on cross-sectional variation in investment good price trends. While their estimate is based on cross-country variation in aggregate (or broad industry) factor share and investment good price trends, I exploit differential exposure across goods, within the U.S., to the secular decline in the real national equipment price—which I view as a more plausible source of exogenous variation. Oberfield and Raval (2014) pursue a very different approach, aggregating the plant-level elasticity of substitution between capital and labor to all value added created within the manufacturing sector. In their framework, substitution to and across intermediate inputs within the manufacturing sector is indirectly accounted for by demand reallocation, while they miss the value added of manufacturing goods that is created in non-manufacturing industries. In general, their approach relies on various aggregation steps, and demand elasticities, increasing model uncertainty. In contrast, my approach allows for directly estimating the overall capital-labor elasticity in production. Relative to that literature, the focus on good-level as opposed to industry-level (or aggregate) factor shares is new. I clarify that my estimated elasticity reflects not only capital-labor substitution within value added of an industry, but also incorporates outsourcing and substitution across intermediate inputs, and note that ignoring the non-homotheticity in consumer demand may bias estimates that are based on aggregate data. To my knowledge, the particular identification strategy of using differential exposure across goods to the secular decline in equipment prices is also novel. I also discuss the relation to that literature in more detail in Section 6.5.

Third, this paper relates to the literature on structural change. As in recent contributions by Boppart (2014) and Comin, Lashkari and Mestieri (2015), I find that allowing for long-run income effects is crucial in order to understand long-run sectoral reallocation. In contrast to that literature's focus on broad sectors, I consider a much more disaggregated economy, which is necessary in order to capture the magnitude of the income effect in the aggregate labor share evolution. As Comin et al. (2015), I consider a log-linear demand system. However, I find that non-homothetic CES preferences (which give rise to such a demand system) are in general not suitable for highly disaggregated data. They imply a joint restriction on substitution and income elasticities, which does not hold in my data. Thus, I use a log-linear approximation to Engel curves (as, e.g., used by Aguiar and Bils (2015)).

3 Theory

In this section I present the baseline model. The aim is to formalize how various types of economic growth affect the aggregate labor share in a neoclassical environment. If growth is biased towards a particular production factor, relative factor prices change. In turn, the optimal production input mix and factor shares change, depending on how substitutable inputs are. An example of such biased growth is technical change that is specific to (equipment) capital goods. In an economy with multiple sectors producing multiple final goods, biased growth generally affects the relative prices of these final goods. In turn, depending on how price elastic consumer demand is, expenditure shares change, which in turn affects aggregate factor shares. For example, the relative prices of those final goods that are capital-intensive in production fall. If consumer demand is sufficiently price elastic, expenditure shares of capital-intensive goods rise, which depresses the aggregate labor share. Moreover, any form of economic growth increases real income. Unless consumer demand is restricted to be homothetic, rising real income affects expenditure shares. For example, if necessities tend to be capital-intensive and luxuries labor-intensive, then economic growth has a positive impact on the aggregate labor share.

The baseline model is intentionally reduced to the elements that I found to be quantitatively relevant. There is no explicit role for consumer heterogeneity. All markets are competitive; thus, there are no pure profits. I focus on intra-temporal consumption decisions and do not model the inter-temporal consumption-savings choice; i.e., I consider a static model repeatedly, and infer capital stocks (and implicitly investment) from observed factor prices and profit maximization. The economy is closed; there is no role for international trade. I discuss all these aspects in Sections 6 and 7.

3.1 Production

There are multiple sectors $i \in I$. In each of them, a representative firm produces a single final good i. The firm uses the production inputs capital K_{it} and labor L_{it} to produce y_{it} units of output according to a constant-elasticity-of-substitution (CES) production technology:

$$y_{it} = A_{it}\left((1-\alpha_i)^{\frac{1}{\eta}}(A^L_{it}L_{it})^{\frac{\eta-1}{\eta}} + \alpha_i^{\frac{1}{\eta}}(A^K_{it}K_{it})^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}}. \tag{1}$$

The parameter α_i controls factor share levels, and varies across sectors. η is the elasticity of substitution between capital and labor, assumed to be constant across sectors. The technology terms A_{it}, A^K_{it}, and A^L_{it} represent, respectively, the state of factor-neutral, capital-augmenting, and labor-augmenting technology.[8]

The representative firm takes the factor prices of labor W_t and capital R_t as given, and maximizes profits. The labor share of good i is defined as $\theta^L_{it} \equiv \frac{W_t L_{it}}{W_t L_{it} + R_t K_{it}}$, and by profit maximization equal to

$$\theta^L_{it} = (1-\alpha_i)\left(\frac{A^L_{it}\tilde{p}_{it}}{W_t}\right)^{\eta-1}, \tag{2}$$

[8]Even though the model is static, I use time indices to indicate which objects are time-varying, because I will consider this static model repeatedly over time.

where $\tilde{p}_{it}$ is the TFP-neutral price of good i, a weighted average of factor prices in efficiency units:

$$\tilde{p}_{it} = A_{it} p_{it} = \left((1-\alpha_i) \left(\frac{W_t}{A_{it}^L} \right)^{1-\eta} + \alpha_i \left(\frac{R_t}{A_{it}^K} \right)^{1-\eta} \right)^{\frac{1}{1-\eta}}. \tag{3}$$

3.2 Consumer demand

Consumers are endowed with $\bar{K}_t$ units of capital and $\bar{L}_t$ units of labor. Consumer demand is characterized by a (common) compensated substitution elasticity σ_t, and by good-specific income elasticities γ_{it}. Expenditure shares $\omega_{it} \equiv \frac{c_{it} p_{it}}{E_t}$ are exogenously given for some base period $t = \tau$. They change over time according to:

$$d \ln \omega_{it} = (1-\sigma_t) d \ln \frac{p_{it}}{P_t} + (\gamma_{it} - 1) d \ln \frac{E_t}{P_t}, \tag{4}$$

where p_{it} is the price of good i, and $E_t = W_t \bar{L}_t + R_t \bar{K}_t$ is total nominal expenditure. The price deflator is defined as $d \ln P_t \equiv \sum_{i \in I} \omega_{it} d \ln p_{it}$, and the budget constraint imposes $\sum_{i \in I} \omega_i \gamma_{it} = 1$. Equation (4) can be interpreted either as an ad hoc specification of demand, or as a first-order approximation to the demand system implied by some underlying primitive utility function.[9]

3.3 Equilibrium

All factor and goods markets are competitive. It is convenient to choose the wage rate as the numeraire ($W_t = 1$). Then, solving for the (unique) equilibrium reduces to finding the rental rate of capital R_t such that given these factor prices, and given the implied good prices, factor and good markets clear.

Formally, an equilibrium consists of factor prices (W_t, R_t), good prices $(p_{it})_{i \in I}$, consumer demand $(c_{it})_{i \in I}$ and expenditure E_t, final good output $(y_{it})_{i \in I}$, and factor input choices (L_{it}, K_{it}), such that

(i) consumer demand $c_{it} = \frac{\omega_{it} E_t}{p_{it}}$, where ω_{it} is exogenously given for $t = \tau$ and evolves according to (4);

(ii) given good prices $(p_{it})_{i \in I}$ and factor prices (W_t, R_t), final good output $(y_{it})_{i \in I}$ and factor input choices (L_{it}, K_{it}) are consistent with profit maximization subject to (1);

(iii) all final good markets clear,

$$c_{it} = y_{it} \quad \forall \, i \in I; \tag{5}$$

[9] See Appendix D for the case of non-homothetic CES preferences.

(iv) and all factor input markets clear,

$$\bar{L}_t = \sum_{i\in I} L_{it}, \text{ and } \bar{K}_t = \sum_{i\in I} K_{it}. \tag{6}$$

3.4 Comparative statics

The aggregate labor share in this economy is denoted by $\bar{\theta}^L_t$ and given by

$$\bar{\theta}^L_t \equiv \frac{W_t\bar{L}_t}{W_t\bar{L}_t + R_t\bar{K}_t} = \sum_{i\in I}\omega_{it}\theta^L_{it}. \tag{7}$$

We are interested in general equilibrium changes in the aggregate labor share in response to changes in the aggregate fundamentals of this economy. These are the growth rates of labor ($g^L_t \equiv d\ln A^L_t\bar{L}_t$), capital ($g^K_t \equiv d\ln A^K_t\bar{K}_t$), and TFP ($g^A_t \equiv d\ln A_t$). The growth rates of labor and capital refer to efficiency units; i.e., to the sum of physical growth and improvements in factor-augmenting technology. Regarding the latter, note that I only consider changes to the respective common components here. Denote by $g_t \equiv d\ln\frac{E_t}{P_t}$ the overall real growth rate, which can also be expressed as $g_t = g^A_t + \bar{\theta}^L_t g^L_t + (1-\bar{\theta}^L_t)g^K_t$ (as in a one-sector model).

To begin with, I consider how good-level labor shares respond to changes in factor prices (in partial equilibrium). Let $\hat{r}_t \equiv d\ln R_t/A^K_t$ and $\hat{w}_t \equiv d\ln W_t/A^L_t$ denote log changes in factor prices per efficiency unit. The change in the labor share of sector i with respect to an (infinitesimal) change in factor prices is then given by:[10]

$$d\theta^L_{it} = (\eta-1)\theta^L_{it}(1-\theta^L_{it})(\hat{r}_t - \hat{w}_t). \tag{8}$$

The following proposition characterizes the response of the aggregate labor share to various types of aggregate growth.

Proposition 1. *The general equilibrium response of the aggregate labor share $\bar{\theta}^L_t$ with respect to the growth rates of TFP (g^A_t), labor (g^L_t), and capital (g^K_t) is given by*

$$d\bar{\theta}^L_t = \frac{\tilde{\eta}_t - 1}{\tilde{\eta}_t}\bar{\theta}^L_t(1-\bar{\theta}^L_t)(g^L_t - g^K_t) + \frac{g_t}{\tilde{\eta}_t}Cov_t(\gamma_{it},\theta^L_{it}), \tag{9}$$

where $\tilde{\eta}_t$ is a convex combination of η and σ_t, $Cov_t(\gamma_{it},\theta^L_{it})$ refers to the cross-sectional covariance between income elasticities and sectoral labor shares (weighted by expenditure shares), and $g_t = g^A_t + \bar{\theta}^L_t g^L_t + (1-\bar{\theta}^L_t)g^K_t$.

[10] More generally, for any neoclassical production function, the substitution elasticity can be defined as $\eta_{it} \equiv \frac{\partial\ln(L_{it}/K_{it})}{\partial\ln(R_t/W_t)}$. Then

$$\frac{\partial\ln\theta^L_{it}}{\partial\ln(R_t/W_t)} = (1-\theta^L_{it})\frac{1}{\theta^L_{it}(1-\theta^L_{it})}\frac{\partial\theta^L_{it}}{\partial\ln(R_t/W_t)} = (1-\theta^L_{it})\frac{\partial\ln(\theta^L_{it}/(1-\theta^L_{it}))}{\partial\ln(R_t/W_t)} = (1-\theta^L_{it})\frac{\partial\ln(W_tL_{it}/R_tK_{it})}{\partial\ln(R_t/W_t)} = (1-\theta^L_{it})(\eta_{it}-1).$$

The specific production technology that is used in this section restricts η_{it} to be constant across sectors and time. My preferred interpretation is that η represents the appropriate average capital-labor elasticity in the economy.

Proof. See Appendix C.1. □

It is perhaps helpful to re-write equation (9) in terms of relative factor prices as:

$$d\bar{\theta}_t^L = (\tilde{\eta}_t - 1)\bar{\theta}_t^L(1 - \bar{\theta}_t^L)(\hat{r}_t - \hat{w}_t) + g_t Cov_t(\gamma_{it}, \theta_{it}^L). \tag{10}$$

The first term in (10) resembles the partial equilibrium response of good-level labor shares in (8) to changing factor prices, except that the capital-labor elasticity in production η is replaced by $\tilde{\eta}_t$. $\tilde{\eta}_t$ captures both the extent of direct substitution in the production process (towards the cheaper production factor) as well as indirect substitution, or re-allocation of consumer demand, towards sectors that use the relatively cheaper factor more heavily.[11] In the data, the weight on η averages 82%, with little time variation. Therefore, $\tilde{\eta}$ is close to η, and capital-labor substitution depends primarily on technology, not on preferences. The second term in (10) represents consumer demand reallocation towards high income elasticity goods in response to rising real income, holding prices constant (i.e., in partial equilibrium).[12]

In what follows, I illustrate Proposition 1 by presenting three special cases that each isolate one channel.

Example 1 (one-sector model): If there is just one sector, then trivially the composition of consumer demand does not matter. I.e., the covariance term is equal to zero and $\tilde{\eta}_t = \eta$. Then, as is well known from the neoclassical growth model, factor shares are stable if and only if (i) either $\eta = 1$ or (ii) all growth is labor-augmenting (then, $g_t^L = g_t^K$ as capital is the elastic factor). If part of growth is investment-specific, such that $g_t^L < g_t^K$, then the labor share increases if capital and labor are gross complements ($\eta < 1$) and decreases if they are gross substitutes ($\eta > 1$).

Example 2 (price effects in consumer demand only): Assume that there are multiple sectors, and that each sector uses either only labor or only capital. Then, η is irrelevant (more precisely, it is not well-defined). Assume also that consumer demand is derived from homothetic CES preferences with substitution parameter σ. Then, it is as if consumers' preferences are defined over labor and capital services directly. Formally, $\tilde{\eta}_t = \sigma$, and again the covariance term equals zero. The effects of various types of economic growth on the aggregate labor share are exactly as in Example 1, with η replaced by σ.

Example 3 (income effects in consumer demand only): Assume that there are two sectors $i \in \{n, l\}$, each operating a Cobb-Douglas production technology ($\eta = 1$). Sector n produces a basic necessity and sector l a luxury good. Sectoral labor shares θ_n^L and θ_l^L may vary across sectors (but are constant over time by virtue of the production technology). Consumer demand is such that compensated price changes do not change expenditure shares ($\sigma_t = 1$), and income elasticities are given by $\gamma_{nt} < 1 < \gamma_{lt}$ (e.g., generalized Stone-Geary preferences could result in such a pattern

[11] Oberfield and Raval (2014) refer to $\tilde{\eta}_t$ as the aggregate capital-labor elasticity of substitution.

[12] Note that the impact of overall growth g_t on the aggregate labor share is dampened (amplified) in general equilibrium if $\tilde{\eta}_t > 1$ ($\tilde{\eta}_t < 1$). To understand this, consider the case of a positive covariance term: the partial equilibrium shift of consumption towards labor-intensive goods in response to an increase in real income, parametrized by g_t, leads to an increase in the relative wage rate, which further increases the labor share if and only if $\tilde{\eta}_t < 1$.

of elasticities at a point in time). Note that since $\sum_{i=n,l} \omega_{it} = 1$, and also $\sum_{i=n,l} \omega_{it}\gamma_{it} = 1$ from the budget constraint, we can solve for the income elasticity of necessities as a function of the luxury good's expenditure share and income elasticity: $\gamma_{nt} = \frac{1-\omega_{lt}\gamma_{lt}}{1-\omega_{lt}}$. Then the source of growth is irrelevant for the evolution of aggregate factor shares. All that matters is the overall growth rate, and how labor-intensive the luxury good is relative to the necessity:

$$d\bar{\theta}_t^L = g_t Cov_t(\gamma_{it}, \theta_i^L) = g_t \omega_{lt} \left(\gamma_{lt} - 1\right) \left(\theta_l^L - \theta_n^L\right). \tag{11}$$

3.5 Taking stock

We can use the content of Proposition 1 and the basic facts of Figure 1 for a preliminary analysis of historical developments in the post-war U.S. economy. To begin with, observe that the basic homothetic model cannot account for the data irrespective of how substitutable capital and labor are. Since part of technical change has been investment-specific throughout the period, the model predicts either an ever falling, or an ever increasing, or an ever constant labor share.[13]

On the other hand, the non-homothetic model is a promising in light of the positive relation between household income and household labor shares. If sectoral labor shares and income elasticities are positively correlated, then economic growth affects the aggregate labor share positively. As the overall growth rate has been roughly comparable throughout (and the covariance is roughly stable over time), the magnitude of this income effect has been comparable over time. Thus, the non-homothetic model has the potential to explain the evolution of the aggregate labor share if capital and labor are sufficiently substitutable, the factor price of capital declined relative to the wage rate, and this decline accelerated starting in the 1980s—as suggested by the basic facts of Figure 1.

Proposition 1 also sets the stage for the remainder of this paper. Labor shares (θ_{it}^L, $\bar{\theta}_t^L$) as well as overall growth g_t are directly observed. I proceed to estimate η, $(\gamma_i)_{i\in I}$, and σ. As some of the underlying micro data is not available for the full sample period in sufficient quality, I will assume that the key elasticities have been stable over time. Finally, I will use model restrictions and further assumptions to infer the $(g_t^L - g_t^K)$ term from observed time series of relative capital prices and real output growth.[14]

4 Empirics

In this section I introduce the data sources, provide descriptive evidence, and discuss identification as well as estimation of the structural parameters. The starting point are the BEA's Input-Output Tables, which contain the necessary information to construct a panel dataset of labor shares at

[13] If one were to allow for substantial time-variation in the capital-labor elasticity, then even the homothetic model could fit the data. Specifically, the capital-labor elasticity must have been close to one prior to the 1980s, and increasing subsequently.

[14] Note that $(\hat{r}_t - \hat{w}_t)$ in equation (10) is not directly observed either, because the term refers to factor prices in efficiency units.

the level of individual goods. Next, I link these to consumption micro data from the Consumer Expenditure Survey (CEX) in order to compute the labor content of consumption baskets in the cross-section of households, and to estimate income elasticities. I then augment the labor share panel with information on equipment capital intensities of individual goods, sourced from the NBER-CES Manufacturing Industries Database, the BEA's Fixed Asset Tables, and Compustat. These equipment intensities, in combination with aggregate time series of relative capital prices, are the crucial ingredient for estimating the capital-labor elasticity of substitution in production. Finally, I discuss how I calibrate the remaining model parameters, given these estimated elasticities.

4.1 Input-Output Tables

Industries: The Bureau of Economic Analysis' (BEA) Detailed Input-Output (I-O) Tables form the basis of the empirical analysis. They are available every five years; I use the 1982, 1987, ..., 2007 editions.[15] In a given year t, the *Make Table* specifies the (dollar) amount of good $i \in I_t$ produced by industry $j \in J_t$.[16] The *Use Table* specifies the amount of production inputs used by each industry, where inputs are both value added (labor, capital, production taxes and subsidies) as well as intermediate inputs. The Use Tables allow for directly computing industry-level labor shares $\tilde{\theta}^L_{jt}$ as the ratio of labor compensation payments to total value added. A few details have to be taken care of in doing so: taxes and subsidies are allocated to labor and capital proportionally, and the portion of proprietors' income that reflects labor compensation has to be imputed and re-classified.[17]

Goods: In the remainder of this paper, good-level labor shares θ^L_{it} are the object of interest. My approach to the data is that consumer demand is defined over goods (e.g., a car) instead of over value added by industry (e.g., value added in the car industry). For consistency, the production functions then have to be specified by good as well, and not by industry. Herrendorf, Rogerson and Valentinyi (2013) label this approach to the data as the final expenditure approach, in contrast to the value added approach. Intuitively, θ^L_{it} is a weighted average of industry-level labor shares, where the weights are given by the fraction of total value added generated in producing good i that is originating in each industry. For expositional simplicity, assume that each good i is produced only by a single industry i, and each industry produces a single good. 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}, can

[15] Previous editions are not suitable as value added is not broken down in labor compensation and other components. The 2012 edition is not yet available (as of October 2018).

[16] In the language of the I-O Tables, goods and services are called commodities. I will refer to both goods and services simply as goods.

[17] See Appendix A.1 for details.

be found by solving the following linear system:

$$\theta_{it}^L = \beta_{it}\tilde{\theta}_{it}^L + (1-\beta_{it})\sum_{j\in I_t}\Gamma_{ijt}\theta_{jt}^L, \quad \text{for } i \in I_t. \tag{12}$$

Define the $N_t \times N_t$ matrix $\Gamma_t = [\Gamma_{ijt}]_{i\in I_t, j\in I_t}$, let $D(z)$ denote a matrix that has the entries of the vector z on the diagonal and zeroes off the main diagonal, 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 (12) is

$$\theta_t^L = D(\beta_t)\tilde{\theta}_t^L + D(\vec{1}-\beta_t)\Gamma_t\theta_t^L, \tag{13}$$

from which we can solve for the vector of final good labor shares,[18]

$$\theta_t^L = \left[E_{N_t} - D(\vec{1}-\beta_t)\Gamma_t\right]^{-1} D(\beta_t)\tilde{\theta}_t^L. \tag{14}$$

As such, these good-level labor shares reflect the full value chain. I find that this is desirable because, first, they are invariant to mere re-allocation of production tasks across industries without changing the capital-labor mix within tasks: Consider, for example, a car company that outsources certain auxiliary tasks (e.g., janitorial or accounting services) to specialized companies listed in some services industries. For the sake of the argument, assume the car company is representative of the car industry. If these auxiliary tasks are more labor-intensive than the core task of manufacturing cars, which remains in-house, then the industry-level labor share will decline. The good-level labor share will, however, remain constant (given that the capital-labor mix within these auxiliary tasks stays unchanged). Second, and relatedly, changes in these good-level labor shares reflect not only substitution between capital and labor within an industry, but also across intermediate inputs. Continuing with the example of the representative car company, suppose it outsources an auxiliary task that it used to perform with the company-wide average labor intensity. Suppose in addition that this auxiliary task is performed with a lower labor intensity in some specialized auxiliary industry. In this example, the industry-level labor share stays constant. However, the good-level labor share does change. Hence, when mapping the model to the data, the capital-labor elasticity η reflects not only capital-labor substitution within industries, but also across intermediate inputs. In other words, capital-labor substitution along the full value chain.

Aggregate labor share: The economy-wide labor share $\bar{\theta}_t^L$ can be computed as a weighted average of good-level labor shares. The appropriate weights are in general final demand weights

[18]The matrix $A_t^{-1} \equiv \left[E_{N_t} - D(\vec{1}-\beta_t)\Gamma_t\right]^{-1}$ is the Leontief inverse. Formally, since A_t is an M-matrix, all entries of A_t^{-1} are non-negative. 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 $D(\beta_t)^{-1}A_t\vec{1} = D(\beta_t)^{-1}\left[E_{N_t} - D(\vec{1}-\beta_t)\Gamma_t\right]\vec{1} = D(\beta_t)^{-1}\vec{1} - D(\beta_t)^{-1}D(\vec{1}-\beta_t)\vec{1} = D(\beta_t)^{-1}\left(E_{N_t} - D(\vec{1}-\beta_t)\right)\vec{1} = D(\beta_t)^{-1}\left(D(\beta_t)\right)\vec{1} = \vec{1}$, which proves the claim.

$(\omega_{it}^{FD})_{i \in I_t}$, which can be readily computed from the Use Tables:

$$\bar{\theta}_t^L = \sum_{i \in I_t} \omega_{it}^{FD} \theta_{it}^L. \tag{15}$$

Moreover, the Use Tables also break down final demand into personal consumption expenditure (PCE), various types of private fixed investment (PFI), government consumption and investment, as well as net exports. Hence, one can, for example, also compute the labor share of aggregate private consumption as

$$\bar{\theta}_t^{L,PCE} = \sum_{i \in I_t} \omega_{it}^{PCE} \theta_{it}^L. \tag{16}$$

Industry classification: The I-O industry classifications are time-varying. I first compute θ_{it}^L for each $i \in I_t$ and each t, and subsequently map these objects into a common set of goods $i \in I$.[19] In sum, I end up with a panel dataset of good-level labor shares and expenditure shares, comprising of 373 goods and 6 time periods, spanning the period 1982–2007.[20]

4.2 Consumer demand: Income elasticities

4.2.1 Data and descriptive statistics

I use consumption micro data from the U.S. Consumer Expenditure Survey (CEX), covering 1980-2015. The CEX follows individual households for five consecutive quarters, recording nominal amounts spent on various consumption categories. It comprises of two surveys with separate samples: an interview survey, covering up to 95% of personal consumption expenditures, as well as a diary survey. I mainly rely on the interview survey. Aggregating quarterly to annual expenditures yields a repeated cross-section of annual household expenditures on up to 524 consumption categories (UCCs). I drop households with missing income information, less than the full four quarters of expenditure information, and if the household head is younger than 25 or older than 65 years. After sample selection, the data set consists of 91,894 households, about 2,500 per year. For the few expenditure categories that are missing in the interview survey, I impute spending based on information on income and expenditure in the diary survey.[21]

Next, I map CEX spending data into the I-O Tables' industry classification system. The mapping is based on a manual concordance assembled by Levinson and O'Brien (2015). The final dataset contains for each year t and each household h: total expenditure E_{ht}, expenditure weights $(\omega_{iht})_{i \in I}$, as well as a vector of household characteristics Z_{ht} that includes income and other demographic information. I use the reported after-tax household income variable, which includes transfers. Aggregating across households yields aggregate CEX expenditure weights $(\omega_{it}^{CEX})_{i \in I}$. Conceptually, they correspond closely to the PCE-based expenditure weights.

[19] See Appendix A.1 for details.

[20] For the model simulation, I interpolate expenditure shares and labor shares linearly in between census years.

[21] See Appendix A.2 for further details on treatment of the CEX data.

Given a household's expenditure weights, its consumption-induced labor share is straightforward to calculate as a simple weighted average of good-level labor shares:

$$\theta_{ht}^{L,household} = \sum_{i \in I} \omega_{iht} \theta_{it}^{L}. \tag{17}$$

Figure 2: Consumer Expenditure Survey: Household labor shares

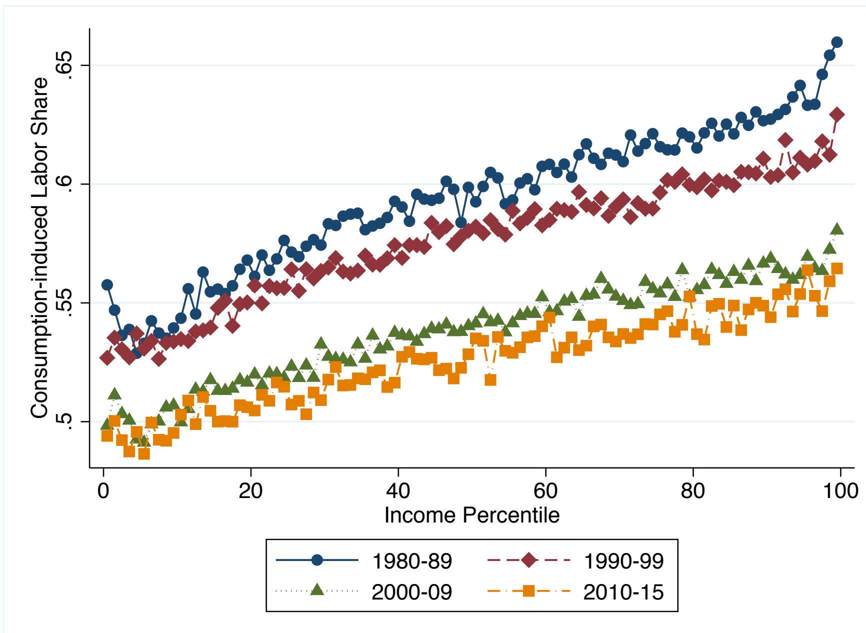

Source: CEX (household consumption by category and income), BEA I-O Tables (labor shares). Income percentiles are defined to be stable over time, so that households in a given income percentile bin, in any year, have the same real income (as a consequence, e.g., there are fewer households in the top income percentile in the 1980s than in the 2010s).

Figure 2 displays averages of these labor shares by income percentile and decade. For each time period, household-level labor shares are positively correlated with household income. I.e., richer households spend more on labor-intensive goods and services as a fraction of their total expenditure. The magnitude is economically significant, as the gap between the 90th and 10th percentile amounts to 6–8 percentage points (for comparison, this is similar in size to the decline in the aggregate labor share). The fact that these labor shares declined over time, conditional on income, reflects primarily technological changes (i.e., θ_{it}^{L} decreased over time for most $i \in I$), and to a small extent substitution towards capital-intensive goods in response to changing prices. The income percentiles in Figure 2 are defined to be constant across years, so that the level of real income on the horizontal axis is constant across years.[22]

[22] Figure 15, reported in the appendix, shows the same statistic with time-varying percentiles.

4.2.2 Identification and estimation of income elasticities

The starting point is the expression for the change in the expenditure share (4) in time t, for good i. To identify the income elasticity, I consider the expression in the cross-section of households h, relative to some reference good $i = 0$:

$$\ln\left(\frac{\omega_{iht}}{\omega_{0ht}}\right) = \zeta_{it} + (1-\sigma_t)\ln\left(\frac{p_{it}}{p_{0t}}\right) + (\gamma_{it} - \gamma_{0t})\ln\left(\frac{E_{ht}}{P_t}\right). \tag{18}$$

Assuming that consumers are facing the same prices, conditional on time and location, prices are absorbed by a good-year fixed effect $\tilde{\zeta}_{it}$. Adding a set of controls Z_{ht} (age, race, family composition, region, urban/ rural) and an error term ξ_{iht}, I estimate

$$\ln\left(\frac{\omega_{iht}}{\omega_{0ht}}\right) = \tilde{\zeta}_{it} + (\gamma_{it} - \gamma_{0t})\ln E_{ht} + \Gamma'_{it} Z_{ht} + \xi_{iht} \tag{19}$$

for all $i \in I \setminus \{0\}$ and t separately in the cross-section of households.

A few comments are in order: First, a standard concern in estimating such a demand system is that measurement error in expenditure on individual goods, as well as the presence of durable goods, implies that total expenditure E_{ht} is correlated with the residual. Therefore, I use current after-tax income, education and occupation (proxies for permanent income) as instruments for total expenditure. The coefficient estimates are generally larger (in absolute value) in the IV specification compared to OLS. Second, the identifying assumptions are that conditional on the observables Z_{ht}, unobserved cross-sectional heterogeneity in prices or preferences is orthogonal to permanent income as proxied by the instruments mentioned above. Third, I estimate relative income elasticities $(\gamma_{it} - \gamma_{0t})$. The levels γ_{it} are easily recovered as their expenditure-weighted average has to equal one (the budget constraint imposes this restriction).[23]

The estimated income elasticities are in fact quite stable over time. A pooled regression of γ_{it} on the time-averages $\bar{\gamma}_i$ yields an R^2 of 0.955.[24] Figure 3 plots time-averages of the estimated income elasticities against labor shares. A complete (aggregated) tabulation is reported in Appendix A.2.3. That appendix also reports the covariance between income elasticities and good-level labor shares for varying degrees of disaggregation. The bottom line is that to capture the income effect quantitatively, one needs to consider a sufficiently disaggregated version of the U.S. economy.

[23] When running this regression separately for each year, one could equivalently let $\ln \omega_{iht}$ be the left-hand side variable in regression (19): the regression coefficient on $\ln E_{ht}$ would then equal $(\gamma_{it} - 1)$. However, when pooling several years, specification (19) is theoretically more attractive since relative income elasticities can be exactly constant over time, while for absolute income elasticities this is generally not the case (because expenditure shares change).

[24] Aguiar and Bils (2015) also use CEX data and a similar log-linear demand specification, with coarser consumption categories, and find that income elasticities are very stable over time.

Figure 3: Income elasticities and labor shares

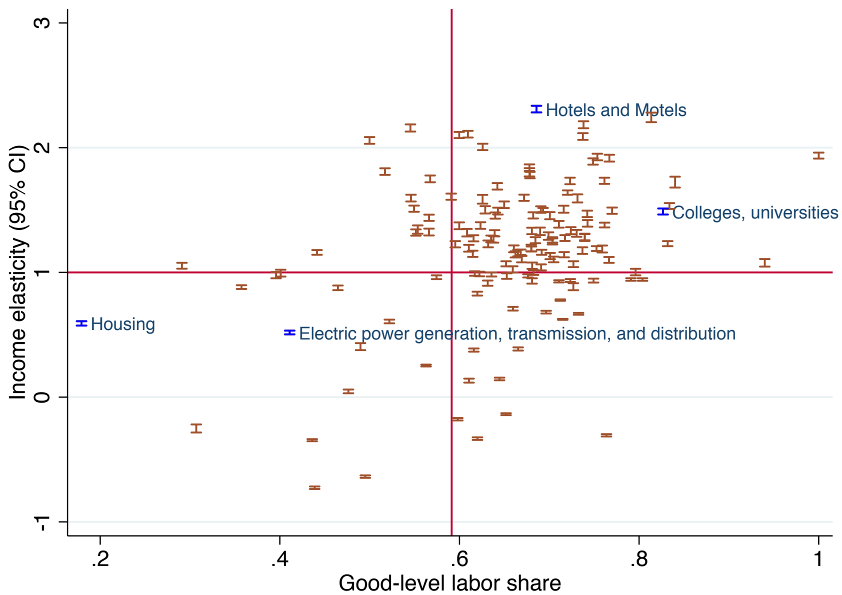

Source: CEX, I-O Tables. This figure displays income elasticities and labor shares, averaged over time. Income elasticities are reported as 95% confidence intervals, corresponding to pooling regression (19) over all sample years. The vertical line indicates the aggregate labor share. A few examples are highlighted.

4.3 Capital-labor elasticity of substitution

Using the FOC for labor (29) and dividing by the one for capital yields:

$$\ln\left(\frac{\theta_{it}^{L}}{1-\theta_{it}^{L}}\right) = \ln\left(\frac{1-\alpha_i}{\alpha_i}\right) + (\eta - 1)\left(r_{it} - w_{it}\right), \tag{20}$$

where r_{it} and w_{it} are log factor prices in efficiency units. The fundamental problem in identifying η is that credible exogenous variation in factor prices is difficult to obtain. My strategy is based on the fact that the relative prices of equipment (including software) and structures capital have evolved very differently over time, as shown in Figure 4. Relative to the price of a consumption good index, the former declined secularly, while the latter remained comparatively flat. I assume that the decline in equipment prices was due to exogenous technical progress. As a consequence, I can treat the fall in capital costs r_{it} that is due to falling equipment prices as exogenous. Crucially, the extent to which this is the case, across goods, depends on how much equipment capital is used in production, relative to structures capital.

4.3.1 Refining the production technology: Two types of capital

Formally, I extend the production technology by adding a second layer. The upper layer remains a CES aggregator of labor L_{it} and capital K_{it}, as defined in (1). I redefine K_{it} to be a composite of

Figure 4: Relative prices of equipment and structures capital

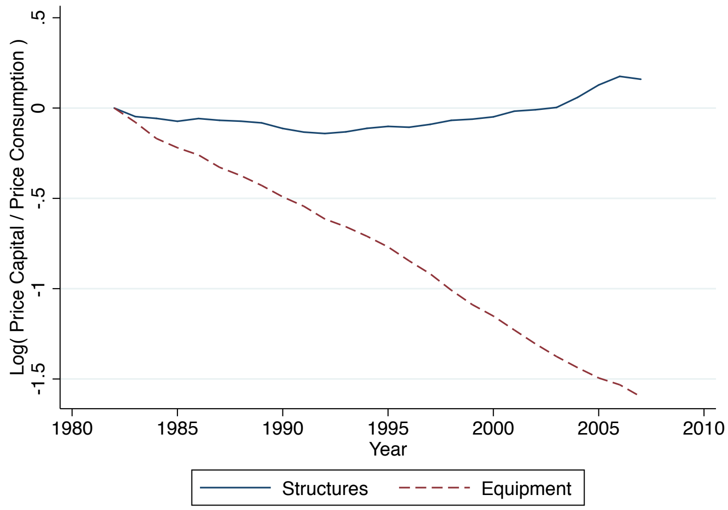

Source: BEA (structures); DiCecio (2009) (equipment & software).

equipment capital K_{it}^{E} and structures capital K_{it}^{S}:

$$K_{it} = \left(\frac{A_{it}^{E} K_{it}^{E}}{\alpha_i^E} \right)^{\alpha_i^E} \left(\frac{A_{it}^{S} K_{it}^{S}}{1-\alpha_i^E} \right)^{1-\alpha_i^E} , \tag{21}$$

where A_{it}^{k}, for $k = E, S$, is capital type k augmenting technology. The representative firm in each sector i chooses the three inputs L_{it}, K_{it}^{E}, and K_{it}^{S} optimally, subject to (1) and (21).

This particular nesting of the three production factors allows for estimating the elasticity between aggregate labor and aggregate capital. The choice of modeling equipment and structures as Cobb-Douglas aggregate is based on data limitations (as discussed below, time-variation in the split of capital income into equipment and structures is difficult to measure; cross-sectional variation is comparatively unambiguous). Moreover, the substitutability between equipment and structures does not affect the evolution of labor shares, at least not to a first order of approximation.

4.3.2 Data: Equipment capital intensities

The I-O Tables only allow for breaking down total value added into labor income and a residual, which I define to be capital income. I split capital income into equipment (which includes software) as well as structures capital income by using data on the nominal stock of equipment $p_t^E K_{jt}^E$ as well as the nominal stock of structures $p_t^S K_{jt}^S$, by industry $j \in J_t$.[25] For the manufacturing sector, the NBER-CES Manufacturing Industry Database (Becker, Gray and Marvakov, 2016) provides

[25]Throughout this paper, I consider a two-way split of capital income into private equipment & 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, and classified them as a separate asset category that also includes software.

this data at the 6-digit industry level. For all other industries, I rely on data from the BEA's Fixed Assets Tables (FAT), which is available at a higher level of aggregation only (62 industries). As an alternative, I use information on public companies from Compustat. Mapping stocks into flows requires further an assumption on the required return on equipment, relative to the one for structures. The required rate of return on a dollar of capital type $k = E, S$, the user cost per dollar of capital, is given by

$$\tilde{R}_t^k = r_t + \delta_t^k - (1 - \delta_t^k)\frac{\mathbb{E}[p_{t+1}^k - p_t^k]}{p_t^k}, \tag{22}$$

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.[26] To compute δ_t^k by capital type and year, I divide current-cost depreciation by current-cost net stock of capital (FAT Table 1.1 and 1.3). For the expected price growth term, I use a five-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.

The equipment intensity $\tilde{\kappa}_{jt}$ of industry j is defined as the ratio of equipment costs to total capital costs:

$$\tilde{\kappa}_{jt} = \frac{\tilde{R}_t^E p_t^E K_{jt}^E}{\sum_{k=E,S} \tilde{R}_t^k p_t^k K_{jt}^k}. \tag{23}$$

A naive strategy would assume that returns on equipment and structures are equal, and thus drop out from equation (23). However, depreciation rates are much larger for equipment; moreover, based on historical experience the expected price decline is also larger.

Given industry j's labor share $\tilde{\theta}_{jt}^L$, I compute the equipment, respectively structures, factor share as

$$\tilde{\theta}_{jt}^E = (1 - \tilde{\theta}_{jt}^L)\tilde{\kappa}_{jt}, \tag{24}$$

$$\tilde{\theta}_{jt}^S = (1 - \tilde{\theta}_{jt}^L)(1 - \tilde{\kappa}_{jt}). \tag{25}$$

Analogously to the computation of good-level labor shares, I map these industry capital shares into good-level equipment (θ_{it}^E) and structures factor shares (θ_{it}^S); similarly, for equipment intensities κ_{it}. Because of data limitations, equipment intensities and factor shares cannot always be computed

[26] 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 five-year moving average of the CPI from both.

separately for every year.[27] However, the time averages $\bar{\theta}_i^E$ and $\bar{\theta}_i^S$, respectively $\bar{\kappa}_i$, suffice for the main empirical exercise. Moreover, using only time-averages for splitting up capital shares has the advantage that estimates of the capital-labor elasticity will not depend on time trends in r_t, which are to some extent ambiguous, as is the choice of the relevant data object for r_t.[28]

4.3.3 Descriptive: Labor shares and equipment shares

Table 1: Exploring the decline in labor shares

	(1)	(2)	(3)	(4)	(5)
t	-0.021***	-0.014***	0.006	0.003	0.002
	(0.005)	(0.005)	(0.009)	(0.009)	(0.014)
$t \times 1_{i \in I_M}$		-0.024**		-0.012	
		(0.010)		(0.010)	
$t \times \bar{\theta}_i^E$			-0.160***	-0.122**	-0.150***
			(0.049)	(0.053)	(0.057)
$t \times \bar{\theta}_i^S$					0.012
					(0.017)

*** p<0.01, ** p<0.05, * p<0.1

Dependent variable: good-level labor shares. 2,172 observations: 362 goods, six time periods (1982, 1987, ..., 2007). Good fixed effects used in all specifications. Standard errors in parentheses (clustered at good level). Unit of t is a decade. Observations are weighted by final demand shares. Labor shares and weights based on I-O Tables (BEA). Equipment and structures intensity based on NBER-CES Manufacturing Database, BEA FAT.

To understand the structural exercise that follows, it is instructive to explore the decline in the aggregate labor share by analyzing time trends in the constructed panel of good-level labor shares. Table 1 displays panel regressions of the good-level labor share θ_{it}^L on good fixed effects and linear time trends. Column 1 reports that the aggregate labor share declined by 2.1 percentage points per decade, or 5.3 points over the period 1982–2007. Column 2 interacts the time trend with a dummy for the manufacturing sector, showing that the decline was stronger for manufacturing goods.[29] Yet, there was still a substantial decline in non-manufacturing (agriculture and services). Column 3 conveys the main point: the entire decline in the labor share has been proportional to the (average) equipment share of goods.[30] Column 4 shows that the stronger trend decline in

[27] In particular, Compustat provides the net equipment capital stock in addition to the total net capital stock only for 1970-1992.

[28] The argument put forward by Karabarbounis and Neiman (2018), that using standard measures of the real interest rate and the user cost formula to back out capital shares leads to unrealistically large swings, also applies to the split between equipment and structures considered here.

[29] Interestingly, Kehrig and Vincent (2017) show that the decline in the labor share of manufacturing industries is largely due to a reallocation of value added towards very productive plants, combined with falling labor shares for these plants. At the same time, surprisingly, the labor share of the median plant increased. This finding is not inconsistent with the capital-labor substitution channel on the good level emphasized in this paper. Instead, it provides guidance for a possible micro-foundation at the firm level.

[30] Using instead the equipment intensity $\bar{\kappa}_i$ produces the same result, regardless of whether the time-average or the 1982 value is used.

manufacturing vanishes when controlling for equipment shares. Finally, column 5 suggests that the structures share is orthogonal to the decline in labor shares.

To summarize, (i) the decline in the U.S. labor share over the previous decades has been concentrated in the production of equipment-intensive goods, and (ii) the quality-adjusted price of equipment capital has exhibited a secular downward trend.

4.3.4 Estimating the capital-labor elasticity of substitution

Let $w_{it} = \ln\left(\frac{W_t}{A^L_{it}}\right)$, $r^E_{it} = \ln\left(\frac{R^E_t}{A^E_{it}}\right)$, and $r^S_{it} = \ln\left(\frac{R^S_t}{A^S_{it}}\right)$ denote factor prices in log efficiency units; let hats denote changes. First, note that the log price of the cost-minimizing good i capital aggregate, r_{it}, changes according to

$$\hat{r}_{it} = \kappa_{it}\hat{r}^E_{it} + (1-\kappa_{it})\hat{r}^S_{it}, \tag{26}$$

where κ_{it} is the equipment intensity of capital. This is a direct application of Shephard's Lemma and does not require any assumptions on the substitutability between equipment and structures in producing good i (neither its value nor whether it is constant). Data limitations require using the average equipment intensity (i.e., $\kappa_{it} = \bar{\kappa}_i$), which is justified in the Cobb-Douglas case, but more generally only valid as a first-order approximation to (26).

I consider two specifications. The first one considers the FOCs for capital and labor (20) in changes. Let $\tilde{\alpha}_i$ denote a good fixed effect, and plug in for the change in the cost of capital $\hat{r}_{it}$ from (26):

$$\ln\left(\frac{\theta^L_{it}}{1-\theta^L_{it}}\right) = \tilde{\alpha}_i + (\eta - 1)\left(\bar{\kappa}_i\hat{r}^E_{it} + (1-\bar{\kappa}_i)\hat{r}^S_{it} - \hat{w}_{it}\right). \tag{27}$$

In words, if say capital and labor are gross substitutes ($\eta > 1$), then the labor share of good i declines if the rental rates of various types of capital fall relative to the wage rate. The sensitivity to these relative rental rates depends on how much of each type of capital is used in production of good i. At this point, it is helpful to orthogonalize (changes in log efficiency unit) factor prices into common components $(\hat{w}_t, \hat{r}^E_t, \hat{r}^S_t)$ and good-specific factor-augmenting technology $(\hat{a}^L_{it}, \hat{a}^E_{it}, \hat{a}^S_{it})$. The latter are mean zero by construction. Then, the estimated equation in the first specification is

$$\ln\left(\frac{\theta^L_{it}}{1-\theta^L_{it}}\right) = \tilde{\alpha}_i + \lambda_t + (\eta - 1)\left(\bar{\kappa}_i\hat{r}^E_t + (1-\bar{\kappa}_i)\hat{r}^S_t\right) + \underbrace{\xi_{it}}_{=(\eta-1)\left(\bar{\kappa}_i(\hat{a}^L_{it}-\hat{a}^E_{it})+(1-\bar{\kappa}_i)(\hat{a}^L_{it}-\hat{a}^S_{it})\right)}, \tag{28}$$

where I added a time fixed effect (λ_t). The time fixed effect absorbs the common component of wage changes, as well as any other common trend that is outside of the baseline model (e.g., uniformly rising markups). I assume that the secular decline in the price of equipment capital was due to exogenous technical progress. Consequently, I use $\bar{\kappa}_i\hat{r}^E_t$ as an instrument for the potentially endogenous regressor $\left(\bar{\kappa}_i\hat{r}^E_t + (1-\bar{\kappa}_i)\hat{r}^S_t\right)$. There are two potential issues: First, changes in good-specific factor-augmenting technology (the error term) might be correlated with average factor shares. This, as well as other potential biases, deserve a separate treatment; I address those in Section 6. Second, factor prices are in efficiency units, which are not directly observed in the data. I assume that changes in capital costs are equal to changes in (quality-adjusted) capital prices, as discussed in 4.3.2. Over the sample period 1982–2007, equipment prices declined at an annual rate of 6.3%. If, say, this number understates the true decline in equipment capital costs, then I will over-estimate the absolute value of $(\eta - 1)$. However, such potential mis-measurement does not affect how much of the aggregate labor share decline is explained by capital-labor substitution: it is only the product of $(\eta-1)$ and the decline in capital costs that matters in that regard. Furthermore, for the estimate $(\hat{\eta} - 1)$ to have the wrong sign, it would have to be the case that the relative cost of equipment capital has *increased* over time. I find that the magnitude of observed price changes makes this improbable.

Measured capital shares sometimes fluctuate close to zero, both because they are measured as a residual, and because they contain profits, which are volatile. In order to be able to take the logarithm, I bound labor shares symmetrically from above and below ($\theta_{it}^L \in [0.05, 0.95]$). In the alternative specification, this is not necessary, as the dependent variable is $\ln \theta_{it}^L$. The alternative specification is based on the firm's FOC for labor (2). Taking the logarithm,

$$\ln \theta_{it}^L = \ln(1 - \alpha_i) + (\eta - 1)\left(\ln(\tilde{p}_{it}) - w_{it}\right), \tag{29}$$

where $\tilde{p}_{it}$ is the TFP-neutral price of good i in t. Using Shephard's Lemma, note that

$$d\ln(\tilde{p}_{it}) = \theta_{it}^L \hat{w}_{it} + (1 - \theta_{it}^L)\hat{r}_{it} = \theta_{it}^L \hat{w}_{it} + (1 - \theta_{it}^L)(\kappa_{it}\hat{r}_{it}^E + (1 - \kappa_{it})\hat{r}_{it}^S). \tag{30}$$

Plugging this expression into (29), rearranging, and adding time fixed effects, the estimated equation in the second specification is:

$$\ln \theta_{it}^L = \tilde{\alpha}_i + \lambda_t + (\eta - 1)\left(\bar{\theta}_i^E(\hat{r}_t^E - \hat{w}_t) + \bar{\theta}_i^S(\hat{r}_t^S - \hat{w}_t)\right) + \underbrace{\xi_{it}}_{=(\eta-1)\left(\bar{\theta}_i^E(\hat{a}_{it}^L - \hat{a}_{it}^E) + \bar{\theta}_i^S(\hat{a}_{it}^L + \hat{a}_{it}^S)\right)}. \tag{31}$$

This equation is very similar to the first specification, except that the dependent variable is $\ln \theta_{it}^L$ instead of $\ln\left(\frac{\theta_{it}^L}{1-\theta_{it}^L}\right)$, and that the relative equipment factor price is interacted with the equipment factor share, instead of the equipment intensity of capital (same for structures). Note that, contrary to the first specification, (31) is based on a first-order approximation to the log labor share of good i. Furthermore, $\hat{w}_t$ is not absorbed in the time fixed effect. I use model generated wage growth (in efficiency units, wages grow at an annual rate of 1.5%). The above argument on potential mis-measurement of capital costs applies to wages as well: if they are mis-measured, $(\hat{\eta} - 1)$ will be biased, but this bias is inconsequential for predicted labor share changes.

Table 2 displays the estimates of $(\eta-1)$. Equipment intensities for non-manufacturing industries are taken from corresponding, more aggregated, industry groups in the BEA's Fixed Asset Table

Table 2: Baseline estimates of $(\eta - 1)$

	Final Demand		Personal Consumption Expenditure	
Dependent Variable	$\log \theta_{it}^L$	$\log \left(\frac{\theta_{it}^L}{1-\theta_{it}^L}\right)$	$\log \theta_{it}^L$	$\log \left(\frac{\theta_{it}^L}{1-\theta_{it}^L}\right)$
	(1)	(2)	(3)	(4)
OLS	0.333**	0.401***	0.250*	0.316***
	(0.134)	(0.153)	(0.128)	(0.116)
IV (only equipment)	0.427***	0.364**	0.328***	0.280**
	(0.133)	(0.158)	(0.127)	(0.122)
N	2,172	2,172	1,407	1,407

*** p<0.01, ** p<0.05, * p<0.1

Each column reports both OLS and IV estimates of $(\eta-1)$. Columns (1)-(2) weigh goods by final demand shares, (3)-(4) by personal consumption expenditure (I-O Tables). The sample size is smaller in the latter, as some goods are pure investment goods. Time and good fixed effects are used in all specifications; standard errors in parentheses (clustered at good level). Columns (1) and (3) refer to equation (31), columns (2) and (4) to (28). Equipment intensities for the manufacturing sector are taken from the NBER-CES manufacturing database; for non-manufacturing they are based on the more aggregated BEA's Fixed Asset Table 3.1 (current-cost net stock of private equipment and software, respectively structures, by industry).

(FAT). (For manufacturing industries, the NBER-CES database provides estimates at the desired detailed level.) The columns titled 'Final Demand' use time-averaged final demand shares as regression weights, the ones titled 'Personal Consumption Expenditure' restrict the regression to consumption goods with the associated weights. OLS estimates for η range from 1.25 to 1.40, depending on specification and sample, and IV estimates from 1.28 to 1.43. As an alternative to relying on more aggregated FAT data, I calculate equipment intensities based on balance sheets of public companies reported in Compustat. The results in Table 3 are somewhat more volatile, depending on sample and specification: OLS estimates range from 1.25 to 1.67, IV estimates from 1.37 to 1.61. All specifications result in statistically significant positive estimates, implying a capital-labor elasticity above one. I use 1.35 as my preferred value for the model analysis to follow, which corresponds to the mid-point of estimates.

4.4 Remaining model parameters

In the previous sections, I discussed how I estimate income elasticities and the capital-labor elasticity of substitution. Both sets of estimates are based on cross-sectional variation. I turn to calibrating the remaining model parameters, given these estimates. The initial model period refers to the year 1982. I repeatedly solve for the static equilibrium over 1982–2007, calibrating growth in factor stocks and technology to match time series of aggregate output and relative factor prices. Intuitively, how much capital prices have been falling is informative for how much of aggregate growth can be attributed to capital. Then, there is a growth residual, which has to be attributed to labor and/ or TFP. Model-generated changes in good prices, driven by differential factor shares and changes in relative factor prices, pin down the consumers' elasticity of substitution

Table 3: Alternative estimates of $(\eta - 1)$: Compustat

	Final Demand		Personal Consumption Expenditure	
Dependent Variable	$\log\theta_{it}^L$	$\log\left(\frac{\theta_{it}^L}{1-\theta_{it}^L}\right)$	$\log\theta_{it}^L$	$\log\left(\frac{\theta_{it}^L}{1-\theta_{it}^L}\right)$
	(1)	(2)	(3)	(4)
OLS	0.282**	0.615**	0.250**	0.672**
	(0.121)	(0.256)	(0.114)	(0.324)
IV (only equipment)	0.407***	0.572**	0.368***	0.609*
	(0.132)	(0.268)	(0.125)	(0.337)
N	2,172	2,172	1,407	1,407

*** p<0.01, ** p<0.05, * p<0.1

See notes for Table 2. The only difference is that for non-manufacturing industries, equipment intensities are calculated based on Compustat firm-level data.

σ.[31] Subsequently, assuming that elasticities have been stable (because the micro data is not available for the earlier time period), the same analysis is extended to 1950–1982.

4.4.1 Base year 1982

In the base year $\tau = 1982$, all expenditure shares and factor shares are matched by construction. It is convenient to normalize the units of goods and factor inputs such that all prices, as well as technology terms $A_{i\tau}$ and $A_{i\tau}^f$, for $f \in \{L, E, S\}$, are equal to one. Then, I set the factor share parameters α_i and α_i^E such that they replicate factor shares $(\theta_{i\tau}^L, \theta_{i\tau}^E, \theta_{i\tau}^S)$ for all goods. By setting factor endowments $(\bar{L}_\tau, \bar{K}_\tau^E, \bar{K}_\tau^S)$ to equal aggregate factor shares $(\bar{\theta}_\tau^L, \bar{\theta}_\tau^E, \bar{\theta}_\tau^S)$, all markets clear as desired at normalized prices.

4.4.2 Sample period 1982–2007

The calibration targets three time series over the sample period 1982–2007: real per capita GDP, as well as the relative prices of equipment capital and structures.[32] The user cost of capital per dollar of capital is assumed to be constant; i.e., growth in rental rates equals growth in the price of each capital type. To match these two capital price time series, I pick the evolution of the capital stock $(\bar{K}_t^E, \bar{K}_t^S)_{t=1983}^{2007}$, which should be interpreted as being in efficiency units. In terms of the model, this is isomorphic to changing (the common component of) capital-augmenting technology. In terms of the mapping to the data, this assumes that all capital-biased technical change is embodied in capital, as measured by relative quality-adjusted prices; that there is no additional, disembodied capital-biased technological progress.

The residual part of economic growth could be attributed to TFP (A_{it}) or labor-augmenting technological progress (A_{it}^L). I impose that the respective good-specific terms are constant; i.e.,

[31] Alternatively, in Section 7.3 I utilize data on final good price changes.

[32] I use the Fisher chained-price index to compute real output changes in the model, consistent with the data.

technology improvements are common to all goods ($\Delta A_{it} = \Delta A_t$ and $\Delta A^L_{it} = \Delta A^L_t$). In the benchmark calibration, I assume that all residual growth is labor-augmenting in the tradition of the neoclassical growth model; i.e., $\Delta A_t = 0$, and ΔA^L_t is chosen to match output growth (jointly with capital growth). Alternatively, I consider tying ΔA^L_t to an index of human capital.[33] I refer to this as the human capital (HC) calibration. Since this index is constructed as years to schooling times returns to education, and there can be other forms of labor-augmenting technological progress, the HC calibration should be interpreted as a lower bound for the loading on A^L_t. I attribute the remaining growth residual to TFP (ΔA_t).

The only remaining parameter to estimate is the consumers' substitution elasticity σ (discussed in the following section). Capital stocks, technology terms, and σ are jointly calibrated to match the data targets.

4.4.3 Consumers' substitution elasticity σ

In the baseline, I use model-generated final good prices (p_{it}) to estimate the elasticity of substitution σ in consumer demand.[34] Since technical progress is assumed to be uniform across goods, variation in the evolution of relative prices is driven by variation in factor shares. In particular, equipment-intensive goods become relatively cheaper over time.

To identify the substitution elasticity σ_t, I have to assume that demand shifters, as well as the effect of controls, are time-invariant: $\zeta_{it} = \zeta_i$ and $\Gamma_{it} = \Gamma_i$. Then,

$$\Delta \ln \omega^{CEX}_{it} = (1-\sigma_t)\Delta \ln \left(\frac{p_{it}}{P_t}\right) + (\gamma_{it}-1)\,\Delta \ln \left(\frac{E_t}{P_t}\right) + \Gamma'_i\,(\Delta Z_t) + \xi_{it}. \tag{32}$$

Then, given prices and a choice of price index, I estimate σ_t by regressing the change in residual aggregate expenditure shares, denoted by $\hat{\omega}_{it}$, on the change in relative prices. $\Delta\hat{\omega}_{it}$ is defined as the change in aggregate expenditure shares net of income effects (and possibly changing aggregate demographics in Z_t):

$$\Delta \ln \hat{\omega}^{CEX}_{it} \equiv \Delta \ln \omega^{CEX}_{it} - (\hat{\gamma}_{it}-1)\,\Delta \ln \left(\frac{E_t}{P_t}\right) - \hat{\Gamma}'_i\,(\Delta Z_t) = (1-\sigma_t)\Delta \ln \left(\frac{p_{it}}{P_t}\right) + \xi_{it}. \tag{33}$$

One could get rid of the price index P_t on the RHS by subtracting, for all goods i, the same equation for some reference good 0. Still, the price index is needed to construct residual expenditure shares as it enters the income effect term. In the sample period 1982–2007, one can use a chain-type index such as the Fisher price index, because demand weights are known for all periods. However, for the earlier out-of-sample period 1950–1982, this is not possible (because only current period demand weights are known). For consistency, I use Laspeyres' index throughout: $\frac{P_{t+1}}{P_t} = \sum_{i\in I} \omega_{it} \frac{p_{it+1}}{p_{it}}$.

[33] Source: Penn World Tables 9.0 (Feenstra, Inklaar and Timmer, 2015), retrieved from FRED (series: HCIYISUSA066NRUG).

[34] Alternatively, in Section 7.3 I use price data in a richer calibration.

In practice, measurement error in prices creates attenuation bias when estimating σ_t. As the model predictions for relative price changes are more reliable in the long run, I prefer using long changes; i.e., changes between t and τ (where τ is either the initial or some interim time period):

$$\Delta_{t,\tau} \ln \hat{\omega}_{it} \equiv \sum_{k=\tau+1}^{t} \Delta \ln \hat{\omega}_{ik} = (1-\sigma_t)\Delta_{t,\tau} \ln \left(\frac{p_{it}}{P_t}\right) + \xi_{it}. \tag{34}$$

When using long changes, the substitution elasticity σ_t is assumed to be constant within each segment. Splitting the sample in three time periods, I cannot reject the null of a constant elasticity.[35] Hence, I assume that σ is constant over time. The point estimate is $\hat{\sigma} = 1.51$ (standard error: 0.07).

4.4.4 Earlier time period: 1950–1981

The calibration of capital stocks and technology is completely analogous to the later period, as the aggregate time series on output and capital prices extend back. I assume that the capital-labor elasticity η and consumer demand parameters have been stable. In particular, I set relative income elasticities $(\gamma_i - \gamma_0)$ to their respective time-averages over the later sample period. As noted earlier, they are very stable over the sample period. Each year, I recover their level by imposing that their expenditure-weighted average equals one, as required by the budget constraint.

5 Quantitative results

In this section I first report on the main result: the evolution of the aggregate labor share in the model economy contrasted with the data. Subsequently, I use the model to decompose changes in the labor share into technological and consumer demand components. Finally, I discuss the failure of the neoclassical model in explaining the behavior of the labor share if preferences are restricted to be homothetic.

5.1 Baseline non-homothetic model results

Figure 5 displays the evolution of the aggregate labor share in the data and in the model.[36] Note that by construction, all series agree in 1982. Let us first focus on the sample period 1982–2007: The model performs well in replicating the overall fall in the labor share, irrespective of which particular calibration strategy is used to attribute residual growth to labor-augmenting and factor-neutral components. The fall in the quality-adjusted relative price of equipment capital was drastic over that time period (Figure 1b), implying that growth was mostly investment-specific. Given that I estimate capital and labor to be gross substitutes ($\eta > 1$), good-level labor shares decrease,

[35] See Appendix B.1 for details.

[36] The data series refers to the labor share constructed using the CEX demand weights as well as the good-level labor shares from the I-O Tables for 1982–2007: $\bar{\theta}_t^L = \sum_{i \in I} \omega_{it}^{CEX} \theta_{it}^L$. Neither CEX expenditure data nor detailed good-level labor shares are available prior to 1982. I use an aggregate time series from the BLS for the U.S. business sector for the earlier time period, re-scaled such that it aligns with the CEX-based one in 1982.

Figure 5: Aggregate labor share in baseline non-homothetic model

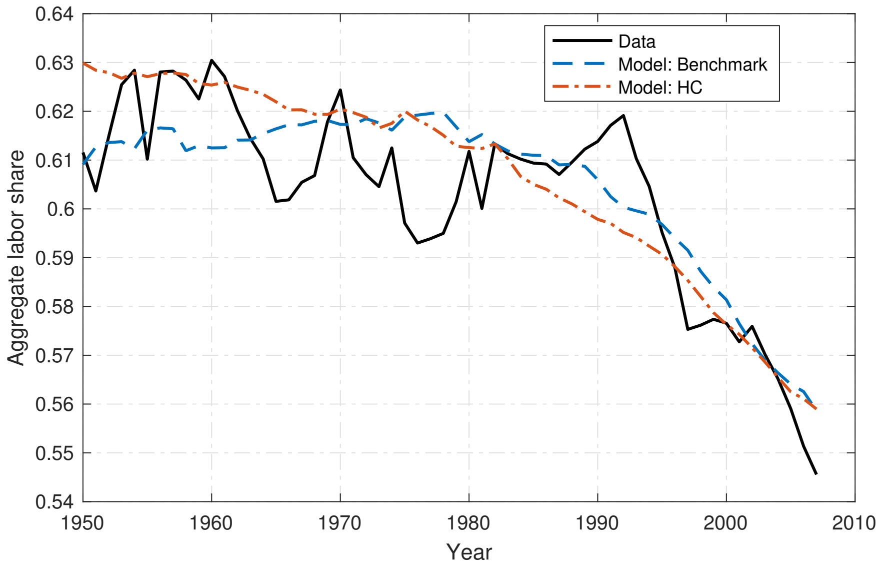

Data sources: BLS, BEA I-O Tables, CEX. In the benchmark model, all residual growth (i.e., that is not accounted for by investment-specific technical change) is labor-augmenting. In the human capital (HC) calibration, growth in labor-augmenting technology is tied to an index of human capital, and the remaining residual is attributed to TFP.

and so does the aggregate labor share. Note that even though I used trends in good-level labor shares over the same time period 1982–2007 to estimate η, the model does not fit the data trend by construction: first, when estimating η I employed time fixed effects to control for economy-wide factors outside of the model; second, changes in consumer demand impact the aggregate labor share as well. For the 1950–1982 period, in the data the aggregate labor share fluctuates without any, or perhaps a slight downward trend. The model does not feature the high-frequency fluctuations, but exhibits a comparable trend behavior.

To understand the differences between the two model calibrations, it is helpful to consider the time series of calibrated technology parameters and capital stocks. The upper two panels of Figure 6 display TFP and the labor-augmenting term A_t^L. The benchmark calibration attributes all residual growth (i.e., that is not accounted for by capital-embodied technological progress) to A_t^L, whereas the alternative only the part that is measured as improved human capital (HC). While there is no cumulative difference from the early 1970s onwards, earlier on there is an unexplained growth residual. Attributing it to labor as in the benchmark has a positive impact on labor shares (given that $\eta > 1$), because it partly offsets the capital bias of growth. In either case the growth rate of the labor-augmenting term has decreased considerably around 1980.

The lower two panels of Figure 6 display the evolution of capital stocks. These are displayed in efficiency units; nominal amounts evolve very differently given drastic changes in relative prices.[37] While the stock of structures grows at a rate slightly below the overall growth rate, equipment grows at a much higher rate.

[37] I compare the model to data on nominal investment and capital stocks in Section 7.4.

Figure 6: Calibrated parameters in baseline non-homothetic model

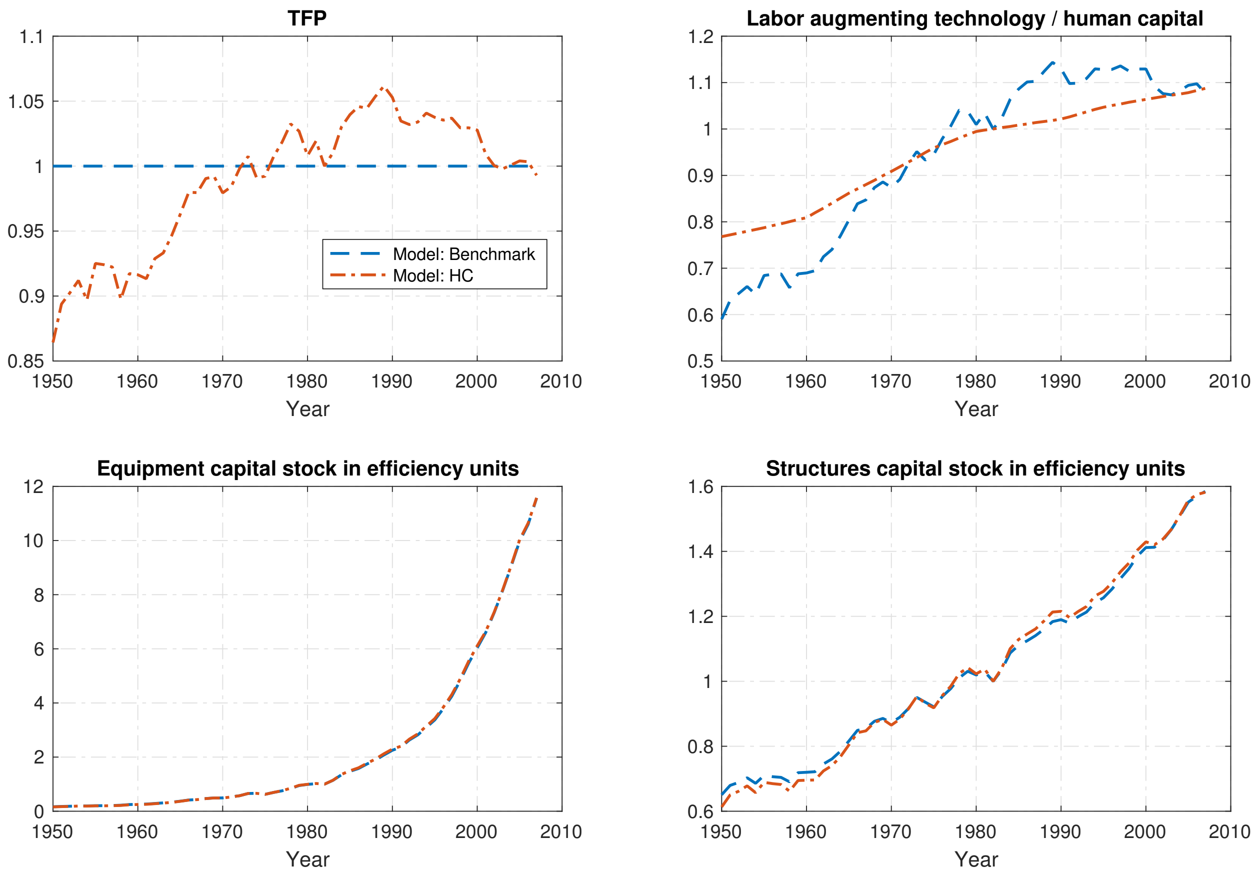

5.2 Decomposing changes in the aggregate labor share

Why was the labor share stable until the early 1980s, and why did it subsequently decline? I provide an answer to this question by using the model to decompose changes in the aggregate labor share into the following three components: (i) capital-labor substitution in the production process, (ii) an income effect operating through non-homothetic preferences, and (iii) a substitution effect on the demand side. I use the benchmark model for this decomposition. To isolate the income effect, I fix the consumers' substitution elasticity σ as in the benchmark, set all income elasticity parameters to one, and re-calibrate technology parameters and capital stocks so that aggregate growth and relative capital prices are unchanged. The difference between this homothetic model and the benchmark is interpreted as income effect. Next, I moreover set $\sigma = 1$ (i.e., demand is characterized by a Cobb-Douglas utility function) and repeat the procedure. The difference between the latter two alternatives is interpreted as substitution effect on the demand side. Finally, capital-labor substitution in the production process is the residual, which is by construction equal to the change in the labor share in the Cobb-Douglas specification.

Table 4 displays the results. Capital-labor substitution in production has been the dominant force, decreasing the labor share by -11.6 percentage points in total. Driven by a steep decline in the price of equipment capital, this force accelerated from the 1980s onwards. On the consumer

Table 4: Decomposition of aggregate labor share changes

	1950–1982	1982–2007	Total
Data	0.2	−6.8	−6.6
Model (benchmark)	0.4	−5.5	−5.0
Technology: K-L substitution	−3.6	−8.0	−11.6
Preferences: Income effect	4.2	3.5	7.7
Preferences: Substitution effect	−0.1	−1.0	−1.1

Data: BLS, BEA I-O Tables, CEX. See text for details on the model.

demand side, rising real income has had a strong positive effect on the labor share throughout (+7.7 points cumulatively), while substitution towards capital-intensive products has played only a minor role (−1.1 points). Looking at the two subperiods, I find that until the early 1980s capital-labor substitution in production was more than offset by the income effect, shifting consumer demand towards labor-intensive goods in proportion to overall economic growth. Subsequently, the former effect dominated.

5.3 Homothetic model results

What does the neoclassical model imply for the evolution of the aggregate labor share if demand is restricted to be homothetic? To answer this question, I first re-estimate the benchmark model under homothetic CES preferences. The estimated consumers' substitution elasticity is lower: $\hat{\sigma} = 0.76$. This is because labor is the factor that is becoming relatively more expensive over time. Thus, when removing income effects, which in the aggregate shift consumption towards labor-intensive goods, this shift is (misleadingly) attributed to substitution effects in the form of a lower estimate for σ. The top left panel of Figure 7 compares the aggregate labor share in that alternative homothetic model to the data, and to the baseline model that allows for non-homotheticities. Without income effects, the decline of the labor share is even more striking. The model fails in particular in the earlier time period, up to the 1980s.

It either model version, it is possible to perfectly match the data on capital prices, real output growth, and the labor share, by treating TFP A_t and the labor-augmenting component A_t^L as free parameters. While this procedure results in reasonable low-frequency movements for A_t and A_t^L in the baseline model, this is not at all the case for the homothetic model. The top right panel of Figure 7 displays linear time trends for TFP in both model versions. The homothetic model requires TFP to fall by around one third over 1950–1982; expressed in annual rates, by more than 1.2% p.a. The intuition is that labor costs have to fall as much as capital costs, but the latter are falling relative to the price of consumption. This is only possibly if TFP is declining. Of course, this is just another way of quantifying the distance between the data and the homothetic model.

What if the estimates of η and/ or σ are wrong, is there some combination of (η, σ) that allows for explaining the behavior of the aggregate labor share without technological regress under homothetic

Figure 7: Non-homothetic vs. homothetic demand

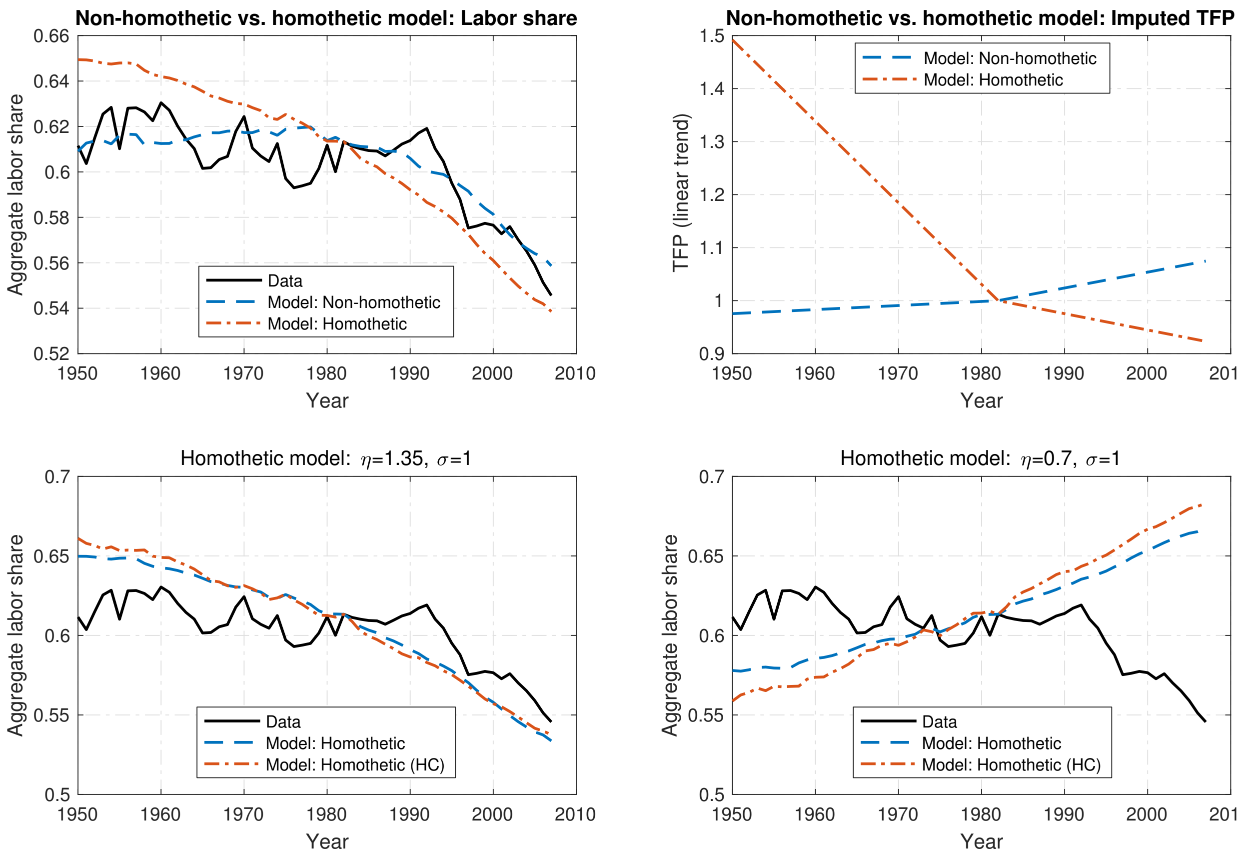

demand? The answer is no, as argued previously. The bottom panels of Figure 7 illustrate this argument for the case of Cobb-Douglas preferences ($\sigma = 1$), with $\eta = 1.35$ as estimated, respectively $\eta = 0.7$. Without technological regress, the model implies an ever declining (increasing) labor share if $\eta > 1$ ($\eta < 1$).

6 Details and context on the capital-labor elasticity η

In this section, I first discuss potential biases that arise in estimating the crucial capital-labor elasticity of substitution η. These are markups and profits, good-factor-specific technical progress, and heterogeneous substitution elasticities. To facilitate the exposition, I analyze the estimating equation (28) in long changes:

$$\Delta \ln \left(\frac{\theta_i^L}{1 - \theta_i^L} \right) = constant + (\eta - 1)\bar{\kappa}_i(\hat{r}^E - \hat{r}^S) + \Delta\xi_i. \tag{35}$$

Subsequently, I put the estimate of η in context by relating it to elasticity parameters in an industry-level framework with explicit input-output linkages, and by discussing the relation to previous estimates in the literature.

6.1 Markups and profits

So far, I assumed that markets are competitive, and thus calculated the capital share θ_i^K as one minus the labor share. Note that if there are instead pure profits, the observed labor share θ_i^L does no longer equal the true labor share of *cost*, but it is still a correct measure of the labor share of *income*. There is no such distinction if markets are competitive. In this section, let θ_i^L denote the labor income share and $\tilde{\theta}_i^L$ the (unobserved) cost share. Let μ_i denote the markup over average cost, so that $(\theta_i^L + \theta_i^K)\mu_i = 1$ and $\tilde{\theta}_i^L = \mu_i \theta_i^L$. The profit share π_i is a function of the average markup: $\pi_i = 1 - 1/\mu_i$.

There is no need to take a stance on the underlying source of market power.[38] The only assumption I maintain is that factor input markets remain competitive. In other words, input decisions are not distorted. Then, (35) still applies, albeit with the ratio of cost shares $\Delta \ln \left(\frac{\mu_i \theta_i^L}{1-\mu_i \theta_i^L} \right)$ on the left-hand side—which we do not observe. Consider first-order approximations to this object around the aggregate markup $\bar{\mu}$ and the aggregate labor share $\bar{\theta}^L$:

$$\begin{aligned} \Delta \ln \left(\frac{\mu_i \theta_i^L}{1-\mu_i \theta_i^L} \right) &\approx \frac{(\Delta \mu_i)\bar{\theta}^L + \bar{\mu}(\Delta \theta_i^L)}{\bar{\mu}\bar{\theta}^L} + \frac{(\Delta \mu_i)\bar{\theta}^L + \bar{\mu}(\Delta \theta_i^L)}{1-\bar{\mu}\bar{\theta}^L} \\ &= \frac{(\Delta \mu_i)\bar{\theta}^L + \bar{\mu}(\Delta \theta_i^L)}{\bar{\mu}\bar{\theta}^L(1-\bar{\mu}\bar{\theta}^L)} \\ &= \frac{\hat{\mu}_i + \hat{\theta}_i^L}{1-\bar{\mu}\bar{\theta}^L}, \end{aligned} \tag{36}$$

where hats denote growth rates.

Thus, if we had data on markups (or cost shares), we could estimate the true elasticity η_0 as

$$(\eta_0 - 1) = \frac{\text{Cov}\left[\bar{\kappa}_i(\hat{r}^E - \hat{r}^S), \frac{\hat{\mu}_i + \hat{\theta}_i^L}{1-\bar{\mu}\bar{\theta}^L}\right]}{\mathbb{V}\left[\bar{\kappa}_i(\hat{r}^E - \hat{r}^S)\right]} = \frac{1}{(\hat{r}^E - \hat{r}^S)(1-\bar{\mu}\bar{\theta}^L)} \frac{\text{Cov}\left[\bar{\kappa}_i, \hat{\mu}_i + \hat{\theta}_i^L\right]}{\mathbb{V}\left[\bar{\kappa}_i\right]}. \tag{37}$$

Instead, in my previous analysis I implicitly set $\bar{\mu} = 1$ and $\hat{\mu}_i = 0$, and estimate

$$(\hat{\eta} - 1) = \frac{1}{(\hat{r}^E - \hat{r}^S)(1-\bar{\theta}^L)} \frac{\text{Cov}\left[\bar{\kappa}_i, \hat{\theta}_i^L\right]}{\mathbb{V}\left[\bar{\kappa}_i\right]}. \tag{38}$$

Thus, the relative bias can be written as

$$\frac{(\hat{\eta} - 1)}{(\eta_0 - 1)} = \frac{(1-\bar{\mu}\bar{\theta}^L)}{(1-\bar{\theta}^L)} \frac{\text{Cov}\left[\bar{\kappa}_i, \hat{\theta}_i^L\right]}{\text{Cov}\left[\bar{\kappa}_i, \hat{\mu}_i\right] + \text{Cov}\left[\bar{\kappa}_i, \hat{\theta}_i^L\right]}. \tag{39}$$

The first term says that if there are pure profits in the economy (on aggregate), so that $\bar{\mu} > 1$, then the estimate has the correct sign but is biased towards one. For example, if $\bar{\mu} = 1.05$, so that

[38] One simple example would be monopolistic competition within each sector, such that markups are a function of the price elasticity of demand across varieties within a sector, which may vary across sectors and over time.

the profit share equals $\bar{\pi} = 1 - \frac{1}{1.05} \approx 4.8\%$, then an estimated $\hat{\eta} = 1.35$ translates into a true substitution elasticity $\eta_0 \approx 1.39$. The intuition is that the observed change in the labor income share actually represents a more drastic change in the labor cost share. As the variation in capital prices is fixed, the model requires a larger substitution elasticity to reconcile the data (i.e., the deviation from the Cobb-Douglas benchmark of $\eta = 1$ must be larger).[39]

The second term says that if markups increased over time in equipment-intensive sectors, relative to the rest of the economy, then $(\hat{\eta} - 1)$ could even have the wrong sign. For a number of available cross-sectional markup estimates for the U.S. economy, I proceed to show that this is not the case.

First, I consider markup estimates from Hall (2018) for 60 KLEMS industries. Hall estimates Lerner indices of market power, which are defined as marginal profit shares (the ratio of price less marginal cost to the price). I map them into the the more disaggregated I-O industry classification system.[40] The Lerner indices are available in the form of linear time trends. Table 5 reports descriptive regressions of labor share time trends on equipment shares, and/or this measure of time-varying markup power (columns 3 and 4). By decade, labor shares fell on average by 2.3 percentage points. While equipment shares account for all of that decline (as demonstrated previously), the Lerner index is only marginally significant, both statistically as well as quantitatively.[41]

Table 5: Labor shares, equipment shares, and markup trends

	(1)	(2)	(3)	(4)	(5)	(6)
Avg. equipment share ($\bar{\theta}_i^E$)		-0.167***		-0.164***		-0.155**
		(0.061)		(0.062)		(0.061)
Hall's Lerner index trend			-0.049*	-0.048		
			(0.027)	(0.029)		
Compustat markup trend					-0.038**	-0.035*
					(0.019)	(0.018)
Constant	-0.023***	0.007	-0.020***	0.009	-0.021***	0.007
	(0.005)	(0.013)	(0.004)	(0.012)	(0.005)	(0.012)

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Dependent variable: good-level labor share trends $\Delta\theta_i^L$. The regression is run in long changes (1982–2007 for labor shares and Compustat markups, and 1988–2015 for the Lerner index obtained from Hall (2018)). 367 observations (goods). Lerner index time trend estimates are available for 60 industries, at a higher level of aggregation. I map them into the I-O classification. Compustat markups computed using sales and cost of goods sold (COGS). Robust standard errors in parentheses. Observations weighted by final demand shares.

Second, I also compute markups for public firms, using Compustat data. In an influential

[39] To understand this, consider an extreme case where the labor share falls from 0.6 to 0.5 in the equipment-intensive sector and stays constant at 0.6 elsewhere. Assume the true profit share equals 0.3, constant across sectors and over time. Then the actual labor-capital cost share ratio in the equipment-intensive sector drops dramatically from $\frac{0.6}{0.1} = 6$ to $\frac{0.5}{0.2} = 2.5$ (and stays constant elsewhere), while I would mistakenly observe a less dramatic drop from $\frac{0.6}{0.4} = 1.5$ to $\frac{0.5}{0.5} = 1$.

[40] An additional detail is that I consider factor shares at the good level, whereas the Lerner indices are estimated for industries. For the regression reported in Table 5, I map industry-level Lerner indices into good-level Lerner indices, in analogy to the computation of labor shares. However, the regression results are very similar when using unadjusted industry estimates instead.

[41] The same pattern holds for labor share trends within the KLEMS dataset, as I show in Table 13 in the Appendix.

paper, De Loecker and Eeckhout (2017) argue that for these firms, markups (and thus variable profit rates) have increased substantially since 1980. I aggregate firm-level data to industry-level, and subsequently good-level, variable profit rates, taking five-year averages around $t = 1982, 1987, ..., 2007$. Variable profit rates are defined as gross profits divided by sales, where gross profits are sales minus cost of goods sold (COGS). As with Hall's markup measure, on a descriptive level the Compustat markup measure explains little of the cross-sectional labor share trend pattern (columns 5 and 6 in Table 5). More formally, when controlling for Compustat markups, the estimated capital-labor elasticity $\hat{\eta}$ is virtually unchanged (Table 6, column 2). This finding continues to hold when adding sales, general and administrative expenses (SG&A) to variable costs (column 3) as proposed by Traina (2018).[42]

Table 6: Controlling for markups when estimating η

	(1)	(2)	(3)
$(\eta - 1)$	0.400***	0.391***	0.400***
	(0.123)	(0.122)	(0.124)
Compustat markup (COGS)		-0.220	
		(0.205)	
Compustat markup (COGS + SG&A)			-0.065
			(0.664)

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Dependent variable: good-level labor shares $\ln \theta^L_{it}$. The regressions are based on equation (31). Columns (2) and (3) add measures of time-varying Compustat markups as controls, computed using sales and either COGS (2) or COGS+SG&A (3). Time and good fixed effects are used in all specifications; standard errors in parentheses (clustered at good level). 2,166 observations (361 goods, 6 time periods). Observations are weighted by final demand shares.

In sum, I conclude that rising markups are not causing an upward bias in my estimate of η. The existence of pure profits, in general, would on the contrary generate a slight downward bias. If profit trends were causing falling labor shares, and if profit trends were correlated with equipment intensities, then my estimation strategy could fail. However, I find that this is not the case, based on the above two independent markup measures.

Several clarifying remarks are in order. First, strictly speaking, my findings do not preclude the possibility that *uniformly* rising markups contributed to the observed fall in the labor share (in the regressions, such an effect would be subsumed by the time fixed effect). If markups did indeed rise across the board and caused the labor share to fall, then the fall in the labor share is over-explained, and there must be other driving forces outside of the model that increased the labor share.

Second, the evidence on rising concentration and rising markups is mixed. While product-

[42] An implicit assumption in these regressions is that the output elasticity of the variable input may be time-varying and sector-specific, but it is not sector-time specific. De Loecker and Eeckhout (2017) provide markup estimates for various assumptions on the production technology; in particular, the headline measure allows for time-sector variation in the output elasticity. They find that the increase in markups is not driven by changes in the output elasticity, but by an increase in the variable profit rate, on which I focus here.

market concentration has increased on the national level, the opposite is the case on the local level (Rossi-Hansberg et al., 2018). The local level is presumably more relevant both for monopsony as well as for monopoly power. Hall (2018) notes that because of substantial sampling variation, the null hypothesis of no change in market power over time cannot be rejected (the p-value for the one-tailed test is 0.11). As for the dramatically increasing aggregate markup reported by De Loecker and Eeckhout (2017), this time-series is sensitive to the choice of aggregation weights (as noted by them, see Figure B.4(b) in the appendix). Their headline measure is constructed as a sales-weighted average. When weighting by cost, the increase post-1980 is muted and comparable to the pre-1980 variation. When relating to objects such as the aggregate labor share, the cost-weighted series is the relevant one.[43]

Finally, depending on the precise definition of markups, in general it is possible that markups are rising without an accompanying fall in the labor share. That would be the case if the wedge between marginal cost and price increases, but this increase does not translate into rising economic profits because it is entirely offset by rising overhead costs—indeed, overhead costs as measured by SG&A in Compustat have increased.

6.2 Good-factor-specific technical progress

Here, I discuss potential biases arising when good-factor-specific technological progress terms are correlated with the equipment capital intensity $\bar{\kappa}_i$, which enters the regressor. Formally, the residual error term in (35) is equal to

$$\Delta\xi_i = (\eta - 1)\left(\bar{\kappa}_i(\hat{a}_i^L - \hat{a}_i^E) + (1 - \bar{\kappa}_i)(\hat{a}_i^L - \hat{a}_i^S)\right). \tag{40}$$

I focus on the equipment capital augmenting term $\hat{a}_i^E$, and set $\hat{a}_i^S = \hat{a}_i^L = 0$. Moreover, to simplify the exposition, I assume $\hat{r}^S = \hat{w} = 0$ for the purposes of this discussion. Then it can be shown that

$$(\hat{\eta} - 1) = (\eta - 1)\left(1 + \text{Corr}\left[\bar{\kappa}_i, \bar{\kappa}_i\hat{a}_i^E\right]\frac{\sigma\left[\bar{\kappa}_i\hat{a}_i^E\right]}{\sigma\left[\bar{\kappa}_i\right]|\hat{r}^E|}\right), \tag{41}$$

where $Corr(x, y)$ denotes the the correlation coefficient between x and y, and $\sigma(\cdot)$ the standard deviation. $\hat{\eta}$ is an unbiased estimate of η if $\bar{\kappa}_i$ is independent of $\hat{a}_i^E$.[44] In general, based on the theory of directed technical change, one would expect that equipment-intensive sectors experienced higher equipment-augmenting technical progress. This would suggest that the correlation term is positive. If this is the case, $(\hat{\eta} - 1)$ does have the correct sign, but is overestimated (given that

[43] To see this, consider the simplest setting. Define the aggregate markup as $\bar{\mu} \equiv \frac{\bar{s}}{\bar{c}}$, where $\bar{s}$ is sales and $\bar{c}$ is cost. Let subscript f denote firm-level objects. Then $\bar{\mu} = \frac{\sum_f s_f}{\sum_f c_f} = \sum_f \frac{c_f}{\sum_{f'} c_{f'}}\frac{s_f}{c_f} = \sum_f \frac{c_f}{\sum_{f'} c_{f'}}\mu_f$; i.e., cost-weights are appropriate. The sales-weighted average increases if the dispersion of markups increases, which is the case in the data. See Edmond et al. (2018) for a more formal argument about welfare, and in particular the striking difference between sales- and cost-weighting reported in Figure 8.

[44] See Appendix B.2.2 for derivations.

$\hat{\eta} > 1$). In terms of the magnitude of the bias, it is important to note that the price of equipment, which is in the denominator of the bias term in (41), fell by more than a factor of four over the sample period 1982–2007 (relative to a consumption good price index, even more so relative to structures). Consequently, systematic differences in $\hat{a}_i^E$ need to be very large in order to affect $\hat{\eta}$ quantitatively.

To assess the quantitative importance of this bias, I perform a Monte Carlo simulation. I assume that equipment-augmenting technology is growing at a constant rate g_i^E, which is good-specific. Table 7 reports the corresponding estimates for equation (31) 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 intensity $\bar{\kappa}_i$. For example, $\sigma(g_i^E) = 0.01$ translates into a standard deviation of 25 log points over 25 years, across goods. For this value and a correlation of 0.5 between g_i^E and $\bar{\kappa}_i$, the corrected estimate of η is equal to 1.333, while I estimated 1.364. In general, the corrected estimate is decreasing in the dispersion of growth rates. As long as this correlation is positive (as suggested by the theory of directed technical change), $(\hat{\eta} - 1)$ has the correct sign, even if the dispersion of growth rates is very large.[45]

Table 7: Sensitivity of $\hat{\eta}$ to equipment-augmenting, good-specific technical progress

$\sigma(g_i^E)$: $\text{Corr}(g_i^E, \bar{\kappa}_i)$:	0.001	0.005	0.010	0.020	0.030	0.040	0.050
-1.000	1.364	1.349	1.278	0.918	0.641	0.623	0.675
-0.500	1.364	1.350	1.315	1.144	1.033	0.966	0.946
-0.250	1.366	1.359	1.276	1.167	1.093	1.071	1.010
0.000	1.366	1.341	1.314	1.227	1.158	1.156	1.103
0.250	1.364	1.349	1.332	1.268	1.199	1.172	1.120
0.500	1.365	1.341	1.333	1.252	1.208	1.180	1.173
1.000	1.363	1.352	1.329	1.279	1.237	1.205	1.179

This table reports, for a given value of $\sigma(g_i^E)$ by column and $\text{Corr}(g_i^E, \bar{\kappa}_i)$ by row, the median IV estimate of $\hat{\eta}$ in equation (31) across simulation runs. In each simulation run, g_i^E is drawn from a normal distribution with mean zero, standard deviation $\sigma(g_i^E)$, and correlation $\text{Corr}(g_i^E, \bar{\kappa}_i)$ with $\bar{\kappa}_i$. Then, the effective change in the log price of equipment capital for good i is correctly computed as $\hat{r}_{it}^E = \hat{r}_t^E - (t - 1982)g_i^E$ (instead of $\hat{r}_{it}^E = \hat{r}_t^E$ as in the baseline).

6.3 Heterogeneous substitution elasticities

The strategy for estimating η requires assuming that this elasticity is homogeneous across goods. This is certainly restrictive. Thus, a relevant question is whether this approach produces an unbiased estimate of the average elasticity $\bar{\eta}$, where the average is taken over potentially heterogenous good-level elasticities η_i for $i \in I$. Suppose the true model is as in (35), but with η

[45] To get a ballpark estimate of the size of this bias: the standard deviation of 5-factor TFP growth in 6-digit manufacturing industries, as reported in the NBER-CES Manufacturing Industry Database, is equal to 0.205 over the sample period 1982–2007 (which is equal to 0.0082 in terms of annual growth rates); the correlation with equipment intensity equals 0.138. Using these two values decreases the IV estimate of η from 1.364 to 1.330.

replaced by η_i. Then it can be shown that

$$(\hat{\eta} - 1) = (\bar{\eta} - 1) + \frac{\text{Cov}\left[\bar{\kappa}_i^2, (\eta_i - 1)\right] - \text{Cov}\left[\bar{\kappa}_i, (\eta_i - 1)\right] \mathbb{E}\left[\bar{\kappa}_i\right]}{\mathbb{V}\left[\bar{\kappa}_i\right]}. \tag{42}$$

As long as η_i is independent of $\bar{\kappa}_i$, the second term in (42) is equal to zero, and $\hat{\eta}$ is an unbiased estimate of $\bar{\eta}$.[46] This orthogonality assumption could be violated if certain goods have a high equipment intensity of capital ($\bar{\kappa}_i$) because capital and labor are very substitutable in production. Note, however, that there is no mechanical correlation in this framework. A higher substitution elasticity η_i leads to a faster falling labor share, but does not directly affect the share of equipment capital income relative to total capital income ($\bar{\kappa}_i$). The latter, more precisely changes in the latter, are regulated by the elasticity of substitution between equipment and structures.

6.4 Explicit input-output linkages

In this paper, I do not model input-output linkages explicitly, but instead specify the production function at the final good level with only labor and (different types of) capital as inputs. Thus, the resulting estimated overall elasticity η represents substitution between capital and labor along the full value chain of a good. In the literature, the production function is often specified at the industry level, meaning that the estimated elasticity reflects capital-labor substitution within an industry only. It can be shown that the overall substitution elasticity η, which I estimate, is a convex combination of the above mentioned within-industry substitution elasticity and two other elasticities that regulate how substitutable intermediate inputs are.[47] Hence, my estimate $\hat{\eta} = 1.35$ is in general consistent with different estimates at the industry level. In particular, if intermediate inputs are highly substitutable, then good-level elasticities are higher than industry-level elasticities.[48]

6.5 Relation to the literature

The literature on estimating the capital-labor elasticity is voluminous. While there is a lot of variation, most estimates are below one. I proceed to discuss the relation to a few key papers.

Karabarbounis and Neiman (2014) is closest to the present paper as their strategy is also based on long-run, cross-sectional variation in the price of capital. While this paper zooms into the U.S. economy and constructs the price of capital by detailed industry (and subsequently for each good) as a cost-share weighted average of national capital prices, Karabarbounis and Neiman (2014) analyze cross-country variation in capital prices and factor shares. Their preferred estimate is 1.25, roughly in line with my findings.

[46] See Appendix B.2.3 for derivations.

[47] Appendix B.2.4 formally establishes this claim.

[48] E.g., see Peter and Ruane (2018), who estimate a long-run plant-level elasticity of substitution between intermediate inputs of 4.3, using data on Indian manufacturing industries.

Antràs (2004) is a prominent example of a paper that estimates the capital-labor elasticity based on time-series data for the U.S. economy. Most specifications, including the preferred one, result in a value significantly below one. The identifying assumption in the baseline specification is that factor-augmenting technological change is constant over time (e.g., $A_t^K = A_0^K e^{\lambda_K t}$). In imposing this functional form assumption, one can control for the unobserved A_t^K and A_t^L variables by simply adding a linear time trend to the FOCs for capital and labor (in logs). To deal with deviations from that linear trend, which can bias the estimate, Antràs (2004) employs a generalized instrumental variables strategy. Government wages, government capital stocks, and population size are used as instruments for prices, respectively quantities (depending on the specification), of capital and labor in the private sector. Since annual deviations from a linear time trend constitute the identifying variation, his estimate can be interpreted as short-run (i.e., annual) capital-labor elasticity. In contrast, my estimate is driven by differential exposure to secular trends in capital prices. Thus, it should be thought of as a medium- to long-run elasticity. That capital and labor are more substitutable over longer time periods is plausible.

In general, the findings in this paper pose a problem for estimates based on aggregate data. The non-homotheticity in demand confounds the estimate of the capital-labor elasticity. Restating Proposition 1 in slightly simplified form:

$$d\bar{\theta}^L = (\tilde{\eta} - 1)\bar{\theta}^L(1 - \bar{\theta}^L)d\ln(R/W) + gCov(\gamma_i, \theta_i^L). \tag{43}$$

Both over time and across countries, it is entirely possible that $\tilde{\eta} > 1$, and even though the researcher has perfect data on $d\ln(R/W)$, she estimates $\tilde{\eta} < 1$. For instance, she could observe falling relative capital costs and increasing labor shares, concluding that $\tilde{\eta} < 1$, while the truth is that the positive income effect dominates the substitution effect. A naive approach to aggregate data simulated from the benchmark model in this paper illustrates this: Using the time series for the aggregate labor and capital shares over 1950–1982 and the true evolution of relative capital costs, I estimate $\tilde{\eta} = 0.98$. The intuition is that the aggregate labor share remained roughly stable, while capital costs fell, which is only possible if the capital-labor elasticity is close to one (provided that the non-homotheticity is ignored).

Oberfield and Raval (2014) estimate a capital-labor elasticity of 0.7 for value added within the U.S. manufacturing sector, significantly below one. Their innovative approach departs from the plant-level value added elasticity, which they subsequently aggregate to the manufacturing sector by estimating auxiliary elasticities of substitution in demand. My estimate refers to a different object because I consider the full U.S. economy, including all input-output linkages across sectors. Thus, I consider in addition not only non-manufacturing goods and services, but also the value added of manufacturing goods that is generated in non-manufacturing industries (e.g., janitors or accountants working for contractors in non-manufacturing industries). In addition, my estimate may differ because it is based on differential exposure to equipment price trends across goods, while Oberfield and Raval (2014) estimate the plant-level value added elasticity based on residual wage variation across MSAs. They run the regression in levels in the cross-section. Since plants located in high-wage regions tend to have higher labor shares, this results in an estimate below one. If residual wages are higher because of unobserved skills (or local factor-augmenting productivity differences), the estimator is biased, and the true elasticity can possibly be larger than one. Oberfield and Raval (2014) counter that argument by employing predicted wage changes based on a region's exposure, measured by regional employment shares, to national employment trends in non-manufacturing industries as an instrument for residual wage levels. Using this Bartik-type instrument results in basically the same estimate as OLS. As shown by Goldsmith-Pinkham, Sorkin and Swift (2018), the Bartik instrument is numerically equivalent to using regional industry shares as instruments. Thus, validity of this instrument requires that regional industry shares are uncorrelated with unobserved skills and factor-augmenting technology terms (in levels, not in changes). This is a strong assumption, as for example employment shares of high-skill industries might be high precisely because the level of (unobserved) skills is high in a region. While my estimation strategy is also based on a shift-share design, the panel structure allows me to control for good fixed effects, so that this issue is mitigated.

In theory, abstracting from identification arguments and hypothetically extending it to the full U.S. economy, the approach proposed by Oberfield and Raval (2014) results in a macro elasticity that coincides with the object I defined as $\tilde{\eta}$ (i.e., the convex combination of the production elasticity η and the consumers' substitution elasticity σ). One advantage is that it allows for using local variation in factor prices. On the other hand, in light of the increasing importance of intermediate inputs, it is heavily dependent on correctly specifying demand. Using plant-level data, the primitive elasticity reflects only capital-labor substitution within a plant, excluding the factor content of intermediate inputs. Various aggregation steps are required to account for final and intermediate input demand reallocation, increasing model uncertainty. In contrast, the upside of my approach is that the estimated primitive elasticity already reflects capital-labor substitution along the full value chain.

7 Model extensions and robustness

This section contains extensions of the baseline model and assesses the robustness of results along a number of dimensions: rising consumer heterogeneity, international trade, final good prices, and investment.

7.1 Increasing income and consumption inequality

The model abstracts from changes in consumer heterogeneity, focusing instead on mean income growth. In this section, I argue that incorporating the rising inequality channel would strengthen income effects, but only to a negligible extent.

Under non-homothetic demand, the distribution of income, or more precisely total consumption expenditures, generally matters for the composition of aggregate demand. Moreover, consumption inequality has increased over the past few decades. However, focusing on the application in this paper, that households' consumption-induced labor shares are monotonically increasing in income does not necessarily imply that more consumption inequality increases the aggregate labor share (i.e., not even in partial equilibrium). Instead, whether this implication is true depends on the relative extent of non-homotheticities in different parts of the consumption distribution. Intuitively, more inequality makes the rich control a larger share of aggregate demand and in addition makes them spend even more on luxuries as a share of their total consumption; however, at the same time the poor will spend even less on luxuries as a share of their total expenditure. Indeed, if demand is derived from generalized Stone-Geary preferences and all consumers are above the subsistence level, then mean-preserving changes to the income distribution have no effect at all on aggregate demand, as these opposing effects exactly cancel out. In agreement with the recent literature on structural change, I find support for income effects that are not leveling off as quickly as Stone-Geary type preferences would predict.[49]

Figure 8: Log-linearity of household labor shares

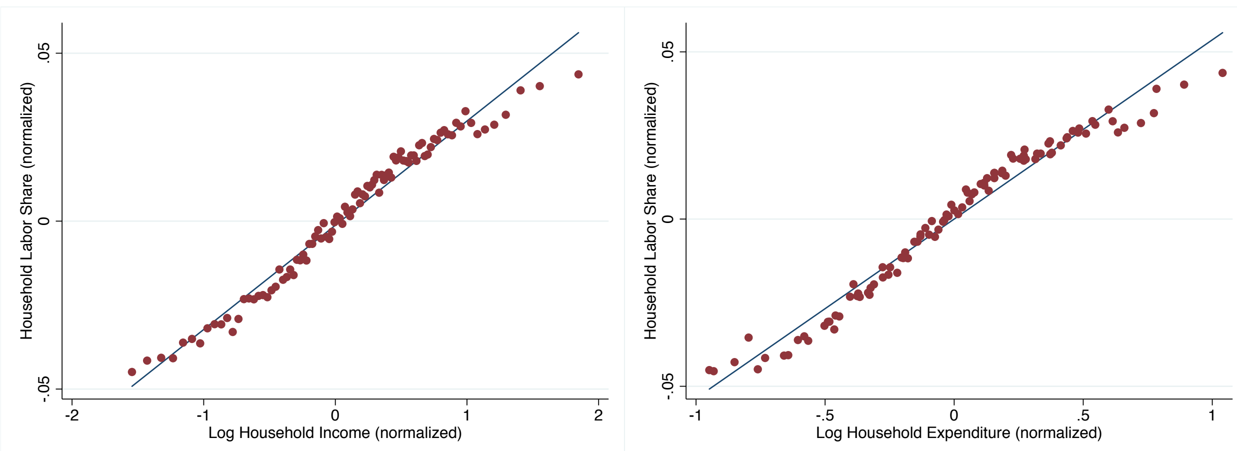

(a) Income (b) Consumption

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).

Figure 8 plots households' consumption induced 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

$$\theta_h^L = c_0 + c_1 \log(E_h), \tag{44}$$

where θ_h^L is the labor share of household h, E_h its total expenditure, and c_0, c_1 are positive scalars.

[49]See for example Buera and Kaboski (2009), Boppart (2014), and Comin et al. (2015).

Let $\log(E_h) \sim N(\mu - \frac{\sigma^2}{2}, \sigma^2)$, so that $\mathbb{E}[E_h] = \exp(\mu)$. Then it can be shown (see Appendix C.2) that

$$\bar{\theta}^L \equiv \frac{\mathbb{E}[E_h \theta_h^L]}{\mathbb{E}[E_h]} = c_0 + c_1 \left(\mu + \frac{\sigma^2}{2} \right). \tag{45}$$

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.[50] Using these numbers, factoring in rising inequality can increase the strength of income effects by $\frac{\Delta\sigma^2/2}{\Delta\mu} = 5 - 14\%$. Since linearity seems to break down at the very top according to Figure 8, I view these numbers as upper bounds. I conclude that while this channel would strengthen income effects, it is small, especially relative to the uncertainty concerning the extent of increasing inequality and the exact shape of the household expenditure–labor share relation.

7.2 International trade

So far, I have abstracted from international trade and assumed that all production is domestic. In this section I show that this simplification does not affect the main findings of this paper: First, accounting for imported goods along the value chain does not change the relation between a household's income and the labor share of its consumption basket. Second, the pattern of factor shares across imports and exports does not mechanically drive the evolution of the aggregate labor share. Third, import exposure does not confound the estimate of the capital-labor elasticity.

The I-O Tables allow for decomposing a dollar of expenditure on good i in year t into payments to domestic labor, payments to domestic capital, and payments to imports.[51] Again, these shares refer to all value added created along the full value chain; i.e., the import share accounts both for imported final and intermediate goods. Given this decomposition of value added, it is straightforward to compute a household's consumption-induced import share as expenditure-weighted average of good-level import shares. I do not find a strong relationship between household income and household import shares (see Figure 16 in the appendix).[52] Figure 9 reports 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. While the second one obviously 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.

Second, if the U.S. imported labor-intensive goods and exported capital-intensive goods, as Heckscher-Ohlin trade theory would suggest, then increased openness to trade would decrease the

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

[51]See Appendix A.3 for details.

[52]This finding is in agreement with Borusyak and Jaravel (2018), who in addition consider consumer packaged goods and automobiles in greater detail.

Figure 9: Household labor shares: Correcting for imports

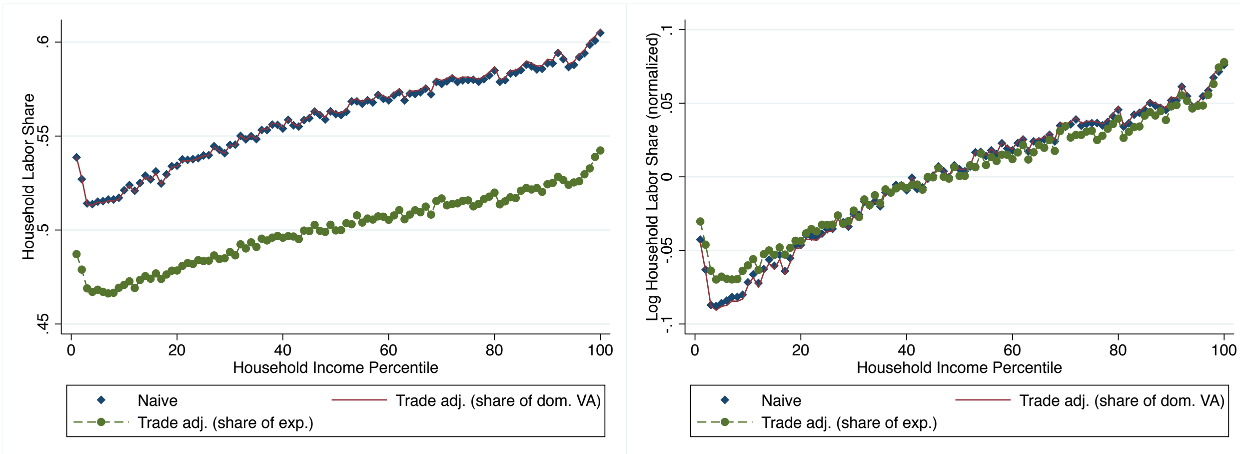

(a) Labor share

(b) Normalized log labor share

Source: CEX, I-O Tables. Households grouped in percentile-year bins according to after-tax household income. Averages over all years (1980-2015) are reported. 'Naive' refers to baseline labor share, which assumes all goods are produced domestically. Trade adjustment factors in imports along the full value chain, such that a dollar of household spending is decomposed into domestic labor payments, domestic capital payments, and imports. Adjusted labor shares are reported as a share of domestic value added and as a share of expenditure.

U.S. labor share. This, however, is not the case. Instead, the factor shares of the baskets of exports and imports are very similar (see Figure 13 in the appendix).[53] Note that the latter is a hypothetical concept that assumes imports are produced with U.S. technology.[54] In that sense, rising import penetration has not directly affected the evolution of the U.S. labor share.

Lastly, exposure to import competition could directly (by outsourcing of labor intensive tasks) or indirectly (e.g., by weakening the bargaining position of domestic workers) decrease labor shares. Potentially, this could lead to a biased estimate of the capital-labor elasticity η, if goods with relatively faster falling capital prices (i.e., equipment-intensive goods) were primarily affected by this channel. Indeed, not only have high-equipment share goods experienced faster falling labor shares (column (1) in Table 8), they also experienced faster growing import shares (4), since they tend to be more tradable. However, the cross-sectional relation between labor and import share changes is rather weak (2), and the correlation between labor share changes and equipment shares is largely unchanged when adding import share changes as a control (3).[55] In that sense, increased openness to trade did not create a spurious association between falling capital prices and falling labor shares.

[53] See also Valentinyi and Herrendorf (2008). A stronger version of this finding is known as the Leontief paradox.

[54] Even though hypothetical, the concept is the right one in this context: For simplicity, say the U.S. used to be in autarky, and now exports financial services to and imports apparel from China. Rising imports and exports mechanically decrease the aggregate U.S. labor share if (and only if) apparel is produced with a higher labor share than financial services within the U.S. Whether or not apparel is produced with a higher labor share in China is irrelevant.

[55] Similarly, adding import shares as a control in the main panel regressions based on equations (28) and (31) does not affect the estimated capital-labor elasticity η. See Table 12 in the appendix.

Table 8: Labor shares, equipment shares, and import shares

Dependent variable:	Labor share trend ($\Delta\theta_i^L$)			$\Delta\theta_i^I$
Independent variable:	(1)	(2)	(3)	(4)
Average equipment share ($\bar{\theta}_i^E$)	-0.165***		-0.161***	0.228***
	(0.049)		(0.053)	(0.053)
Import share trend ($\Delta\theta_i^I$)		-0.092	-0.019	
		(0.058)	(0.053)	
Constant	0.006	-0.018***	0.006	-0.007
	(0.009)	(0.005)	(0.009)	(0.008)

*** p<0.01, ** p<0.05, * p<0.1

The regression is run in long changes (1982–2007). 369 observations (goods). Robust standard errors in parentheses. Observations weighted by final demand shares.

7.3 Good prices

In the baseline model, TFP growth is common across goods: $\Delta A_{it} = \Delta A_t$. 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.[56] Introducing good-specific TFP growth allows the model to match the price data perfectly.

However, the main results are hardly affected, as Table 9 reveals. To understand this, 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 1 (and decreases in some other sector), then the consumed quantity of good 1 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. Similarly, the income effect is not affected by good-specific TFP growth as long as real income growth is equal. This leaves merely the substitution effect on the consumer side to (potentially) change. As it turns out, the effect on this channel, which is small to begin with, is minor. I conclude that the results are not sensitive to allowing for a richer pattern of price changes.

7.4 Investment

So far, I treated consumption as the sole component of aggregate demand. In this section I incorporate investment spending for the following two purposes. First, investment is a sizable fraction of aggregate demand, and differs from consumption with respect to its sectoral composition and changes therein over time. Second, in the baseline approach I back out capital stocks from profit-maximizing behavior of firms, given data and assumptions on the user cost of capital. Introducing investment allows for constructing the capital stock as a result of past investment

[56] The data source 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). More detailed data is not available for the time period considered.

Table 9: Aggregate labor share changes: Sensitivity to alternative price data

	1950–1982	1982–2007	Total
Data	0.2	−6.8	−6.6
Model (common TFP growth)	0.4	−5.5	−5.0
Technology: K-L substitution	−3.6	−8.0	−11.6
Preferences: Income effect	4.2	3.5	7.7
Preferences: Substitution effect	−0.1	−1.0	−1.1
Model (good-specific TFP growth)	0.1	−4.9	−4.8
Technology: K-L substitution	−3.6	−8.0	−11.6
Preferences: Income effect	4.0	3.6	7.6
Preferences: Substitution effect	−0.3	−0.5	−0.8

instead, and thus answering the following two related questions: How do the investment rates implied by the baseline model compare to the data? If we assume that in the data investment rates are measured correctly, how different are the model predictions?

7.4.1 Investment as a component of aggregate demand

Regarding the first objective, Figure 10 plots the change in the aggregate labor share for different definitions of final demand. To reiterate, so far I have equated final demand with private consumption. Thus, I have studied the share of labor in the total cost of producing all consumption goods, which we can refer to as the consumption labor share. The aggregate labor share is falling a bit faster, by about one percentage point, when accounting in addition for private investment spending. Government purchases and net exports do not alter the aggregate trend.[57] Over the sample period 1982–2007, the investment labor share fell by 8.9 percentage points, while the one of consumption fell by 4.9 points. Statistically, we can further decompose this difference into a within-sector and a between-sector, or reallocation, component. As before, a sector does not correspond to an industry, but to a final good (or service), reflecting value added at all stages of production. The top right panel of Table 10 reveals that investment spending shifted towards labor-intensive goods to a much smaller extent than consumption. Note that the reallocation component reflects both substitution as well as income effects. For consumption, as shown 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. Moreover, relative to consumption, investment spending is also concentrated in sectors with faster declining labor shares.[58]

As investment spending, relative to consumption, has been fairly flat without any apparent trend (see Figure 14 in the Appendix), I do not model the optimal consumption-savings choice. Instead,

[57] See Figure 13 and the discussion in Appendix A.1 for further details.

[58] 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 (2018) on the implications for modeling structural change.

Figure 10: Change in labor share for various definitions of final demand

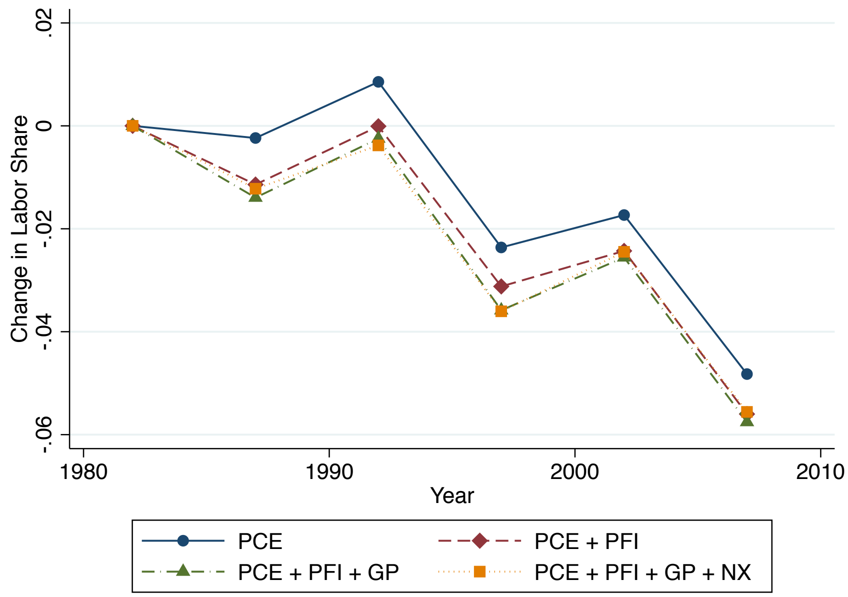

Source: BEA I-O Tables, own computations. Final demand components are personal consumption expenditures (PCE), private fixed investment (PFI), government purchases (GP) and net exports (NX). Reflects all private sectors, government sectors are excluded (GP refers to the government buying goods that are produced in the private sector).

I take nominal investment rates directly from the data. As for the composition, I model investment into equipment & software (E), respectively structures (S), as homothetic CES aggregators:

$$I_t^k = A_t^k \left(\sum_{i \in I} \left(\omega_i^k \right)^{\frac{1}{\sigma^k}} \left(q_{it}^k \right)^{\frac{\sigma^k - 1}{\sigma^k}} \right)^{\frac{\sigma^k}{\sigma^k - 1}}, \quad k \in \{E, S\}, \tag{46}$$

where q_{it}^k is the quantity of good i used for type k investment in period t. The I-O Tables allow for a breakdown of total private fixed investment into the two types for the sample period 1982–2007, for each final good $i \in I$. The level parameters ω_i^k are chosen such that expenditure shares in the model agree with the data in the base year 1982. I estimate σ^E, similarly to the substitution elasticity σ in consumption, from variation in equipment investment shares in response to price changes over time.[59] As in Section 7.3, I utilize price data from the BEA Industry Accounts.[60] For structures, I impose $\sigma^S = 1$ (i.e., expenditure weights are constant over time), as the construction sector accounts for more than 80% of structures investment spending throughout.

[59]Since all prices are normalized to one in the base year 1982 (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 $k = E, S$. For the substitution elasticity, I find that $\hat{\sigma}^E = 1.29$ (standard error: 0.045).

[60]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 equipment, relative to consumption, would not decline as fast as in the targeted data series (DiCecio, 2009).

Table 10: Change in labor share: Aggregate consumption vs. investment

	1950–1982			1982–2007		
	Between	Within	Total	Between	Within	Total
Data						
Total (PCE+PFI)				2.0	−7.9	−5.8
Consumption (PCE)				2.4	−7.3	−4.9
Investment (PFI)				0.9	−9.7	−8.8
Model						
Total	3.6	−3.6	0.1	2.2	−7.6	−5.4
Consumption	4.3	−3.4	0.9	2.9	−7.1	−4.2
Investment	0.1	−4.2	−4.1	0.1	−9.2	−9.0

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 $f \in \{Total, PCE, PFI\}$, between component computed as $\sum_{t=2,\ldots,T} \sum_{i \in I_t} (\omega^f_{i,t} - \omega^f_{i,t-1})\theta^L_{i,t-1}$; within as residual: $\sum_{t=2,\ldots,T} \sum_{i \in I_t} \omega^f_{i,t}(\theta^L_{i,t} - \theta^L_{i,t-1})$.

The bottom panels of Table 10 report the results of this exercise.[61] Looking at the sample period 1982–2007, 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, labor shares are falling faster in those sectors because of the steep decline in the price of equipment capital (given that $\eta > 1$).

7.4.2 Consistency between investment and capital stocks

In the baseline approach followed in this paper, I essentially back out capital stocks from the first-order conditions of profit maximizing firms, feeding in an exogenous factor price of capital:

$$p^k_t \tilde{R}^k_t = p_{i,t} \frac{A_{i,t} \partial F_i(K^E_{i,t}, K^S_{i,t}, A^L_{i,t} L_{i,t})}{\partial K^k_{i,t}}, \quad k \in \{E, S\}, \tag{47}$$

where p^k_t is the relative price of type k capital and $\tilde{R}^k_t$ the required return (per dollar of capital). I take p^k_t from the data and assume $\tilde{R}^k_t$ is constant over time, which is a conservative assumption as applying the user cost formula (22) yields a declining required return post-1980s because of the declining real interest rate.[62] In this baseline approach, the constant $\tilde{R}^k$ is a normalization: By virtue of the calibration, the model fits capital shares in the base year $\tau = 1982$ perfectly.[63] A

[61] 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 aggregate data.

[62] Moreover, even when using moving averages of the realized real interest rate, this leads to unrealistically volatile factor shares.

[63] The baseline model is not truly dynamic, but a repeated static model. Hence, it is possible to choose 1982 as the base year even though the model analysis spans 1950-2007.

higher $\tilde{R}^k$, for $k = E, S$, decreases the imputed real as well as nominal capital-output ratios (and increases the calibrated production function share parameter α_i given that $\eta > 1$), but does not affect growth rates of factor shares.

Given the standard law of motion of capital, the baseline model implies time series of nominal investment rates that can be compared to the data.[64] Let the stocks of equipment and structures evolve according to

$$K_{t+1}^k = (1 - \delta_t^k) K_t^k + I_t^k, \quad k \in \{E, S\}, \tag{48}$$

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.[65] 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 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 $(K_t^k)_{t=1950}^{2007}$, since

$$\bar{\theta}_t^k = \frac{p_t^k \tilde{R}^k K_t^k}{Y_t}, \quad k \in \{E, S\}, \tag{49}$$

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 asset price relative to consumption). In turn, the law of motion (48) implies real investment I_t^k, which in turn can be translated into nominal investment rates i_t^k:

$$i_t^k = \frac{p_t^k I_t^k}{Y_t}, \quad k \in \{E, S\}. \tag{50}$$

To sum up, the baseline approach followed in this paper implies time series of nominal investment rates i_t^k through the equations (49), (48), (50); 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.[66] Given this choice of required return, the left panels of Figure 11 compare investment rates implied by the baseline model to the data. The content of these graphs concerns the later period 1982–2007. While in the data the overall investment rate stays roughly constant at on average 21%, in the baseline model the implied investment rate increases to an average of 26%. The baseline model implies missing

[64] Note that in this section, the 'baseline model' refers specifically to the one introduced in the previous Section 7.4.1, with consumption expenditure shares matching aggregate (PCE) data, and investment (PFI) as a component of final demand. It makes more sense to compare this model to the data on investment rates than the main one discussed in the earlier sections, which is reflecting CEX spending data and does not feature investment expenditures.

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

[66] I use time-consistent aggregate BEA data for nominal investment and consumption shares.

investment of on average 5.1% of output.[67]

An immediate follow-up question is: if we assume that the nominal investment data is true and let the capital stock evolve accordingly, how different are capital stocks and factor shares in such an alternative model? I refer to this as the I-Model, to contrast it to the baseline. In the I-Model, nominal investment rates i_t^k and capital prices p_t^k are taken from the data. For capital market clearing, rental rates of capital adjust; i.e., $\tilde{R}_t^k$ is endogenous (and thus time-varying). I impose that the I-Model replicates the baseline model equilibrium in the initial year 1950 (in particular, $\tilde{R}_{1950}^k = \tilde{R}^{k,\star}$). While other model parameters are borrowed from the baseline model, I re-calibrate 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).

The right panels of Figure 11 compare nominal capital-output ratios in the baseline to the ones generated by the I-Model. The difference can be interpreted as missing capital. Missing total capital is in relative terms lower than missing investment, reflecting in particular the low depreciation rate of structures (which include residential structures), albeit missing equipment capital climbs to 18% of output.[68]

Figure 12 contrasts factor shares of the baseline model to the I-Model. Recall that by construction, the trends over 1950–1982 agree. There are a number of interesting observations: First, as investment is lower in the I-Model post-1982 (relative to the baseline), capital stocks are lower, and consequently (given that $\eta > 1$) capital shares grow at a slower rate. However, even though the nominal capital-output ratio is flat post-1982, capital shares nevertheless increase. This is because the physical capital-output ratio still increases in response to investment-specific technical change. The point is 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 p_t^K the relative price of the investment good. Then the FOC of capital is

$$p_t^K \tilde{R}_t = \alpha^{\frac{1}{\eta}} k_t^{-\frac{1}{\eta}} \tag{51}$$

and the equilibrium capital share is

$$\theta_t^K \equiv \frac{\tilde{R}_t p_t^K K_t}{Y_t} = \alpha^{\frac{1}{\eta}} k_t^{\frac{\eta-1}{\eta}} = \alpha \left(p_t^K \tilde{R}_t \right)^{1-\eta}. \tag{52}$$

The law of motion of capital (48) can be re-written in terms of k_t and the growth rate of physical

[67] The volatility of implied investment rates reflects the absence of adjustment costs. As this paper is not concerned with cyclical fluctuations, I abstract from them.

[68] Note that there are periods, for example the 1960s, where the equipment investment rate goes up in the data, while the equipment capital-output ratio decreases. This is because both objects are reported in nominal terms (indeed, with investment-specific technical change, physical ratios are non-stationary). Equipment investment tends to be high when its price is declining fast, leading to a decreasing nominal capital-output ratio.

Figure 11: Investment rates and capital-output ratios

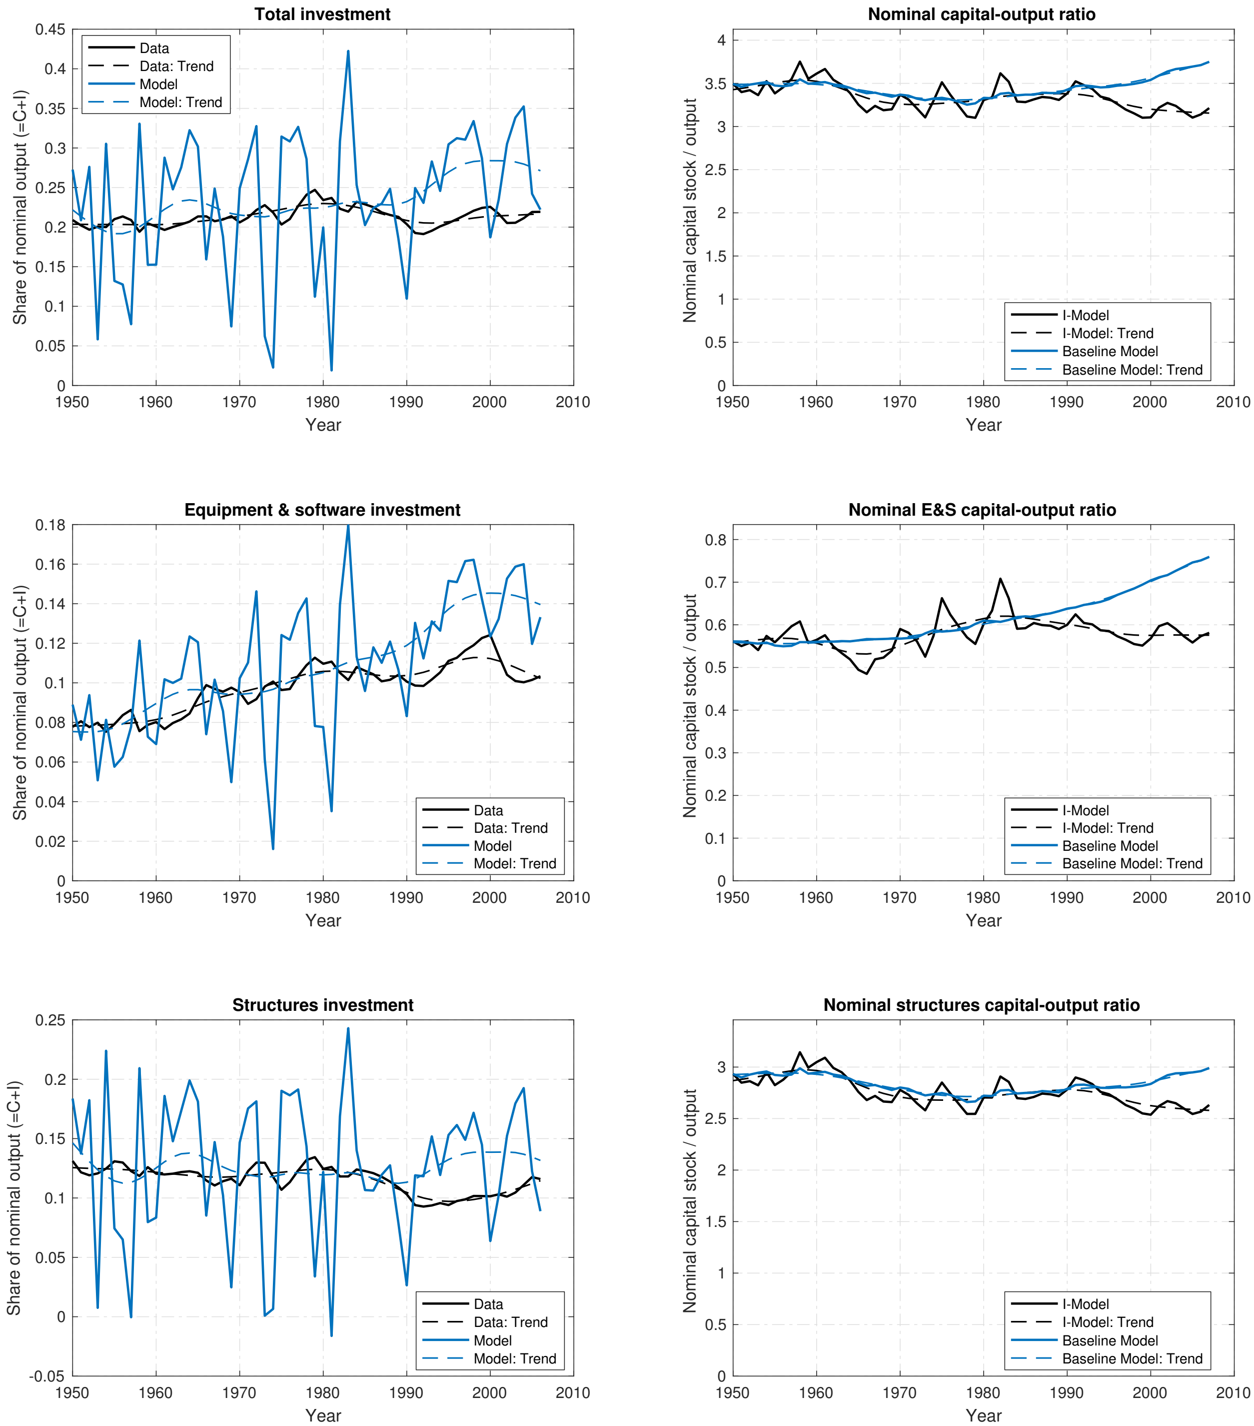

Data: Investment rates are constructed as $\frac{PFI}{PFI+PCE}$, where private fixed investment (PFI) and personal consumption expenditures (PCE) refer to nominal BEA aggregates. The trend lines are obtained from HP-filtering the data with smoothing parameter 100. The I-Model exactly matches the investment data; the capital stock is then constructed accordingly from its law of motion. The baseline model backs out capital stocks from profit-maximizing behavior of firms, given observed capital prices; required investment rates are then computed given the same law of motion of capital.

Figure 12: Factor shares: Different treatment of investment data

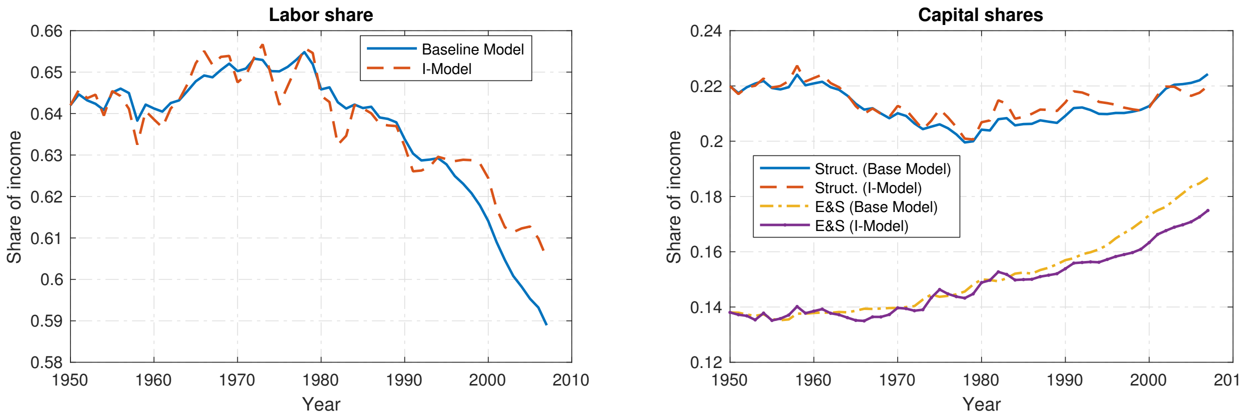

Baseline model includes investment spending as a component of aggregate demand, but does not impose consistency between investment and capital. I-Model in addition imposes consistency between investment and capital stocks, matching the aggregate BEA investment share.

capital, $g^K_{t+1} \equiv \frac{K_{t+1}}{K_t} - 1$, as

$$(g^K_{t+1} + \delta)k_t = \frac{i_t}{p^K_t}. \tag{53}$$

As p^K_t is decreasing at a higher rate post-1982, from (53) it is apparent that the growth rate of physical capital increases given a roughly constant nominal investment rate. This increases k_t, and since $\eta > 1$, from (52) it follows that the capital share increases. The simple one-sector model does of course not feature the counter-acting forces of non-homothetic demand; still, the intuition applies. Quantitatively, the cumulative change in the labor share is -3.7% in the I-Model, as opposed to -5.3% in the baseline.

Second, as the capital share increase is more modest in the I-Model and $\tilde{R}^k$ is constant in the baseline, (52) illustrates why $\tilde{R}^k_t$ has to increase over time in the I-Model for capital market clearing. This requires an increasing wedge between observed real riskfree rates and the rate used by firms when making investment decisions, reflecting the arguments in Caballero, Farhi and Gourinchas (2017), and Karabarbounis and Neiman (2018).

To summarize, the fall in the labor share is somewhat smaller, though still sizable when constructing the capital stock from observed investment data instead of imputing it based on assumptions on the user cost of capital. Investment-specific technical change leads to capital deepening in real terms even if nominal investment rates are constant. Since the capital-labor elasticity is above one, capital deepening is also reflected in increasing capital shares.

8 Concluding remarks

This paper reconciles the time series of the U.S. labor share in the post-war era. Based on a neoclassical framework, I show that its evolution has been driven by two counteracting forces. On the one hand, capital-labor substitution, driven by investment-specific technical change, has put downward pressure on the labor share. On the other hand, an income effect arising from aggregate economic growth has led to a shift of consumption towards more labor-intensive goods and services. The method I follow in this paper is to first estimate the key elasticities using disaggregated production data and consumption micro data, and then use a parsimonious model framework to quantitatively account for the effects of falling equipment prices and aggregate economic growth on factor shares. As observed in the data, abstracting from short-run fluctuations, the model generates an aggregate labor share that is relatively stable until the early 1980s, and declining thereafter. With constant substitution and income elasticities, the model identifies an accelerating bias of technological progress towards equipment capital as the main culprit for the decline in the labor share of national income.

Building on the results of this paper, there are several directions for future work. The first concerns a better understanding of why labor-intensive goods and services tend to be luxuries. One hypothesis is that new goods are more labor-intensive, and high-income households are early adopters. Empirically, investigating this hypothesis requires going beyond the framework in this paper with a fixed set of industries and goods, and to analyze product level data instead.[69]

Second, technological changes are exogenous in this paper. To inform policy, incorporating endogenous and directed technical change might be important. In such a setting, the income distribution can be allowed to feed back into the evolution of factor shares through endogenous improvements in technology, in addition to the consumption channel studied in the present paper.

[69] Scanner data does provide the link between households and products; however, computing labor shares at the detailed product level is challenging.

A Data appendix

A.1 BEA Input-Output Tables

The BEA's I-O Tables, available quinquennially 1982–2007, allow for a breakdown of aggregate final demand into personal consumption expenditures (PCE), private fixed investment (PFI), government consumption and investment (GP), as well as net exports (NX). Given final good labor shares θ^L_{it}, for $i = 1, ..., I_t$, $t = 1982, ..., 2007$, and final demand component expenditure weights ω^f_{it} for $f \in \{PCE, PFI, GP, NX\}$, I compute the labor share of component f as $\bar{\theta}^{L,f}_t = \sum_{i \in I_t} \omega^f_{it} \theta^L_{it}$. Figure 13 reports labor shares by year and final demand component. In particular, note that for imports, this number corresponds to the (hypothetical) share of labor in producing the basket of imported goods in the U.S., using the same technology that is currently used to produce domestic output in these sectors.

Proprietors' income and taxes: I follow Gollin (2002) and in particular Valentinyi and Herrendorf (2008) in employing the economy-wide assumption for splitting the ambiguous part of proprietors' income into payments to labor and capital. I use the BEA's GDP-by-Industry tables, available at the 2-digit NAICS level. Specifically, industry value added is broken down in three parts, as in the more detailed I-O Tables: 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 non-corporate consumption of fixed capital), and into ambiguous income (non-corporate 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 (i.e., labor shares are computed as $\frac{coe_i}{coe_i + gos_i}$ after reallocating part of proprietors' income).

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 1997–2007, are relatively minor. I manually create 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.

Figure 13: Labor share (changes) for aggregate final demand components

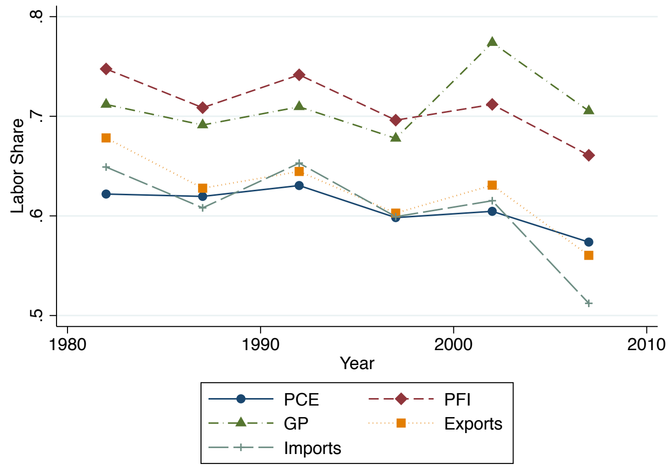

(a) Levels

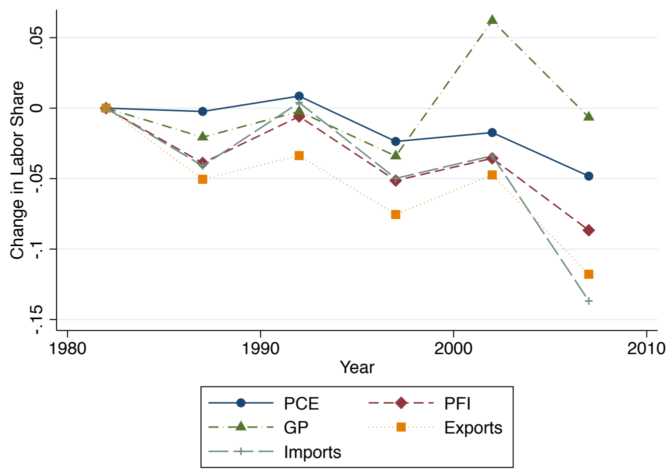

(b) Changes

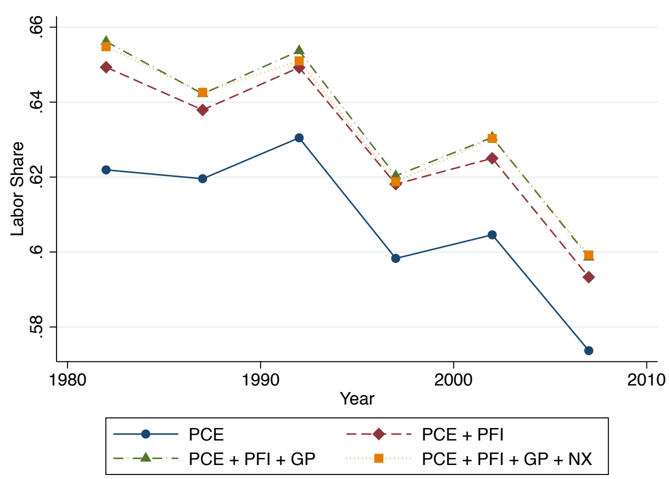

(c) Aggregate Labor Share

Figure 14: Nominal investment ratio

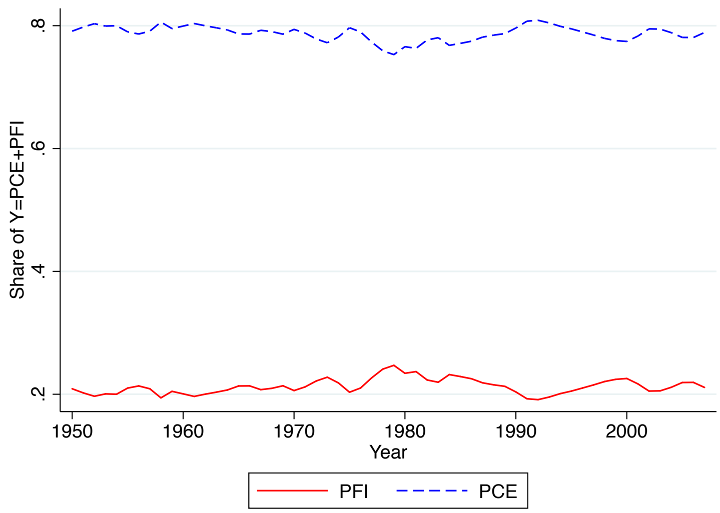

Source: BEA. Nominal annual personal consumption expenditures (PCE) and private fixed investment (PFI).

A.2 CEX

A.2.1 Construction of CEX dataset

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 defined by the BLS itself and captured by the variable *RESPSTAT*. The final sample consists of 91,894 households and spans the time period 1980-2015.

Owner-occupied housing: In general, the CEX measures out-of-pocket expenditures. Hence, while it reports the appropriate rent for renters, for home owners 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 home owners have to be subtracted. Since home owners' rental equivalence is not reported prior to 1982, I impute it based on later survey waves (1982-1990). 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).

Income concept: I use after-tax household income as reported in the CEX (*FINCATAX*) and add the net rental equivalence for home owners to it (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. Specifically, for each of these consumption categories, I regress annual household expenditure as a fraction of household income on log income, demographic controls (see above), and calendar year.

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.2.2 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 (2015). 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 home owners, used car expenses).

Producer vs. 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 in the model consumer demand is specified in producer prices (likewise, all production side data is already in producer prices).

Aggregate consumption in the CEX vs. PCE: A concern with the CEX data is that its representativeness of aggregate consumption (PCE) has been declining over time (see Garner, Janini, Passero, Paszkiewicz and Vendemia (2006)). First, my estimated income elasticities are robust to measurement error as long as measurement error is of the form discussed in Aguiar and Bils (2015); i.e., as long as measurement error is household-specific and / or good-specific, and not household-good-specific. Second, these discrepancies beg 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 weights seem preferable because they are clearly more reliable for aggregate trends. On the other hand, the estimated income elasticities of, e.g., health or education expenditures strictly speaking only correspond to the fraction of these expenditures that are out-of-pocket expenditures, as recorded in the CEX. I choose to use CEX weights for the main results reported in Section 5, and as a robustness check PCE weights in Section 7.4.1 (where I also include aggregate investment demand). I am comforted by the fact that the comparison between model and data is similar in both cases.[70]

A.2.3 Additional CEX results

Figure 15 reports household-level labor shares by income percentile and decade. Contrary to Figure 2 (reported in the main text), income percentiles are time-varying, though the overall picture is similar. Figure 16 displays household-level import shares, which exhibit a slight inverted U-shape in household wealth.

Table 11 summarizes estimated income elasticities, expenditure shares, and labor shares. The reported values are aggregated to the 2-digit I-O level, and averaged over time. Figure 17 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). 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.

[70]Regarding the aggregate labor share trend, the main source of discrepancy is that the level of health expenditures is lower in the CEX. The health sector has a relatively high (and increasing) labor share, and its expenditure share is growing over time.

Figure 15: Household labor shares (time-varying percentiles)

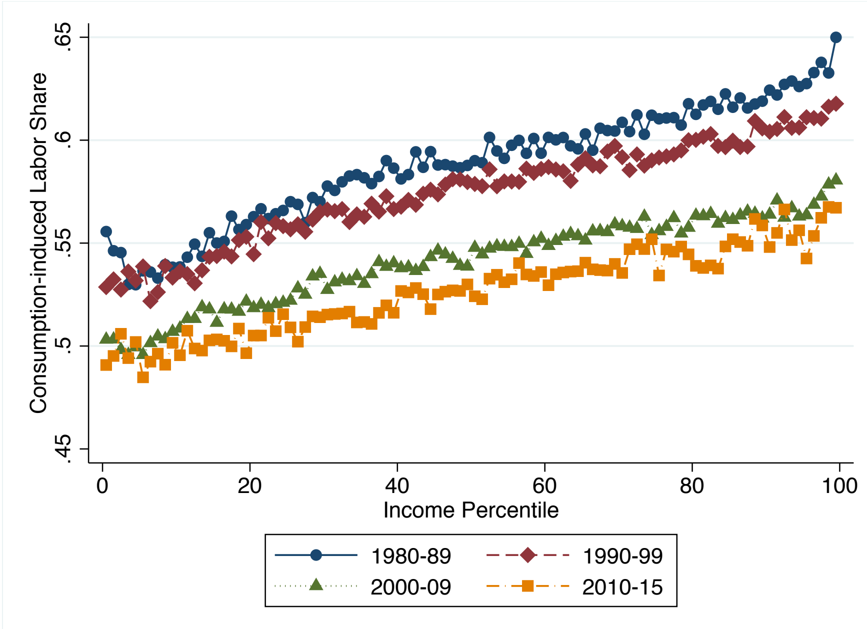

Source: CEX (household consumption by category and income), BEA I-O Tables (labor shares). Income percentiles are time-varying, so that households in a given income percentile bin typically have a higher level of real income in later years.

Figure 16: Household import shares

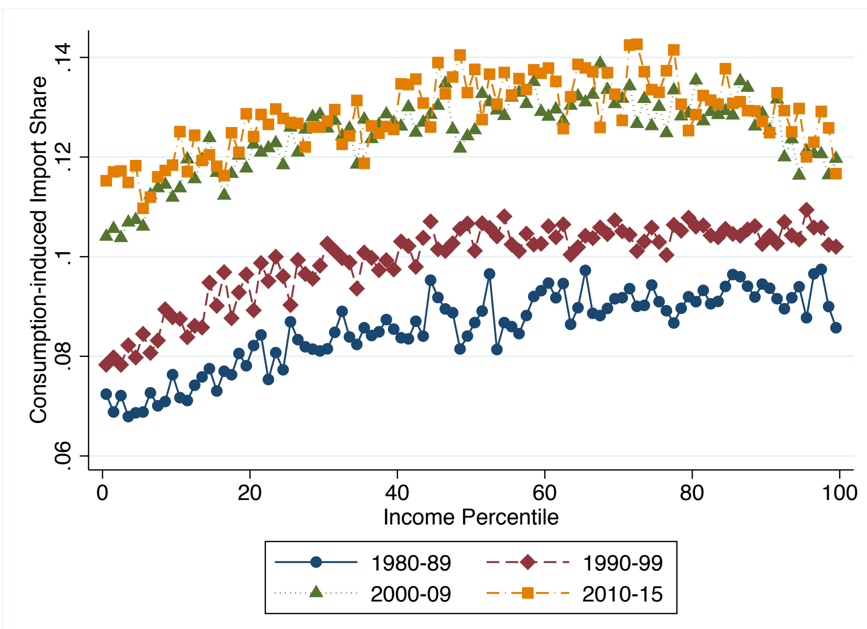

Source: CEX (household consumption by category and income), BEA I-O Tables (labor shares, import shares). Income percentiles are time-varying, so that households in a given income percentile bin typically have a higher level of real income in later years. Latest I-O Table is 2007; thus, data for 2010-15 reflects CEX spending data of that year but import shares of 2007.

Table 11: Expenditure shares, income elasticities, and labor shares, at the 2-digit I-O level

I-O Sector (2-digit)		ω_i^{CEX}	ω_i^{PCE}	γ_i	θ_i^L	$\theta_i^L - \bar{\theta}^L$
11	Agriculture (nurseries, floriculture)	0.001	0.001	1.953	0.517	-0.072
21	Mining	0.002	0.000	1.076	0.297	-0.292
22	Utilities	0.051	0.031	0.747	0.436	-0.153
23	Construction	0.035	0.000	1.739	0.761	0.173
31	Manufacturing I (food, apparel...)	0.110	0.102	0.494	0.572	-0.017
32	Manufacturing II (wood, chemical...)	0.037	0.046	0.452	0.431	-0.158
33	Manufacturing III (cars, electronics...)	0.128	0.060	1.152	0.676	0.088
42	Wholesale trade	0.054	0.034	0.645	0.715	0.126
48	Transportation	0.024	0.026	1.421	0.745	0.156
49	Warehousing, couriers, postal service	0.002	0.001	-0.123	0.757	0.168
4A	Retail trade	0.121	0.115	0.712	0.729	0.140
51	Information	0.042	0.033	0.980	0.505	-0.084
52	Finance and insurance	0.055	0.060	1.770	0.650	0.061
53	Rental and leasing	0.010	0.008	1.564	0.580	-0.009
54	Professional and technical services	0.009	0.016	1.370	0.727	0.139
56	Administrative and waste services	0.005	0.005	1.629	0.718	0.129
61	Educational services	0.019	0.025	1.480	0.795	0.207
62	Health care and social assistance	0.031	0.168	1.406	0.747	0.158
71	Arts, entertainment, and recreation	0.007	0.011	2.070	0.666	0.077
72	Accomodation and food services	0.061	0.060	1.797	0.715	0.127
81	Other services	0.060	0.051	1.425	0.732	0.143
oo	Housing	0.135	0.148	0.451	0.179	-0.410

Source: CEX, I-O Tables. Housing consists of real estate (2002 I-O code 531000) and owner-occupied dwellings (code S00800). Expenditure weights ω_i correspond to final good expenditures, not value added (thus, agriculture has a tiny expenditure share). For various reasons, CEX and aggregate (PCE) expenditure shares differ: e.g., for health care and social assistance, the CEX share is much lower, because it captures only out-of-pocket expenditures. The expenditure-weighted average income elasticity is 1; thus, luxuries are characterized by $\gamma_i > 1$. All reported values are time-averages over 1982–2007.

Figure 17: Cross-sectional covariance of income elasticities and labor shares

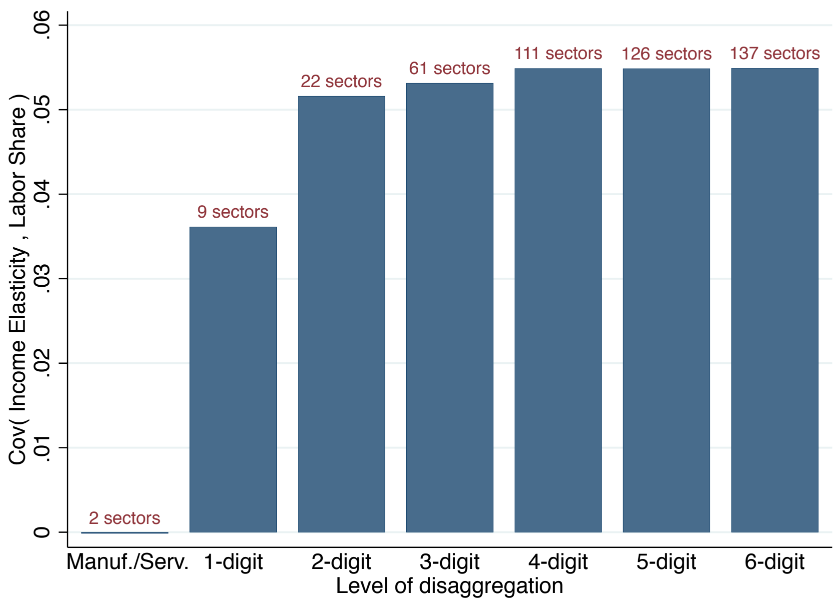

Source: CEX, I-O Tables. This figure displays the (time average over the) cross-sectional covariance of income elasticities and labor shares, for various levels of disaggregation, at the final good level. The rightmost bar corresponds to the level of detail used in the benchmark (Detailed I-O Tables, 6-digit classification, 137 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 (126 goods), ..., all the way to 1-digit sectors (9 goods) and manufacturing vs services (2 goods).

A.3 Accounting for import intensity

The I-O Tables provide data on the use of each good (or service) by final demand category (consumption, investment, government, etc.) and using industry. They also contain the total value of imports, for each good. A naive approach to accounting for imports would just subtract net imports from domestic final demand. This, however, is not just empirically wrong (since the value of imported intermediate goods is substantial), but also infeasible: for many goods, the resulting net domestic final demand would be negative as a result of this procedure. Because the value of imported goods is not available by using industry, the assumption made by the BEA—which I follow—is that each industry uses imports of a given good in the same proportion as the imports-to-domestic supply ratio of that good. E.g., overall 20% of iron and steel mill products were imported in 2002, and the assumption is that each industry that uses iron and steel mill products as an input purchases 80% from domestic producers and 20% from abroad. Imposing this assumption, one can solve for the domestic total requirements matrix, and subsequently for the import share $\theta^I_{it} \in [0, 1]$ of each good (in each year). For the details of these calculations, I refer to Horowitz and Planting (2014).

Table 12 adds the (time-varying) import share θ^I_{it} as a control in the regression that estimates the capital-labor elasticity η. The results are largely unchanged.

Table 12: Controlling for import intensity

I. NBER-CES Manufacturing Database, BEA Fixed Asset Tables:								
	Dependent Var.: $\log \theta^L_{it}$				Dependent Var.: $\log\left(\frac{\theta^L_{it}}{1-\theta^L_{it}}\right)$			
	OLS		IV		OLS		IV	
$(\eta - 1)$	0.333**	0.333**	0.427***	0.434***	0.401***	0.377**	0.364**	0.337**
	(0.134)	(0.137)	(0.133)	(0.138)	(0.153)	(0.153)	(0.158)	(0.159)
$\gamma^{IM-Share}$		-0.002		0.022		-0.177		-0.200
		(0.083)		(0.082)		(0.244)		(0.244)
II.NBER-CES Manufacturing Database, Compustat:								
	Dependent Var.: $\log \theta^L_{it}$				Dependent Var.: $\log\left(\frac{\theta^L_{it}}{1-\theta^L_{it}}\right)$			
	OLS		IV		OLS		IV	
$(\eta - 1)$	0.282**	0.277**	0.407***	0.401***	0.615**	0.585**	0.572**	0.539*
	(0.121)	(0.119)	(0.132)	(0.132)	(0.256)	(0.265)	(0.268)	(0.276)
$\gamma^{IM-Share}$		-0.049		-0.036		-0.205		-0.219
		(0.078)		(0.072)		(0.257)		(0.251)

*** p<0.01, ** p<0.05, * p<0.1

Source: NBER-CES Manufacturing Database, BEA FAT, Compustat, BEA I-O Tables. Time and good fixed effects used for all specifications. $N = 2,172$ (362 goods, 6 time periods). Standard errors (in parentheses) are clustered at the good level. $\gamma^{IM-Share}$ is the coefficient on the import share θ^I_{it}, which is added as a control in this table. See section 4.3.4 for a description of different specifications.

B Additional results

B.1 Time-varying elasticity of substitution σ_t in consumer demand

In the main text, σ_t is restricted to be constant over time. I have considered splitting the sample in three sub-periods. The results are displayed in Figure 18: the estimates are not precise enough to allow for a time-varying σ_t.

Figure 18: Estimates of σ_t for different subperiods

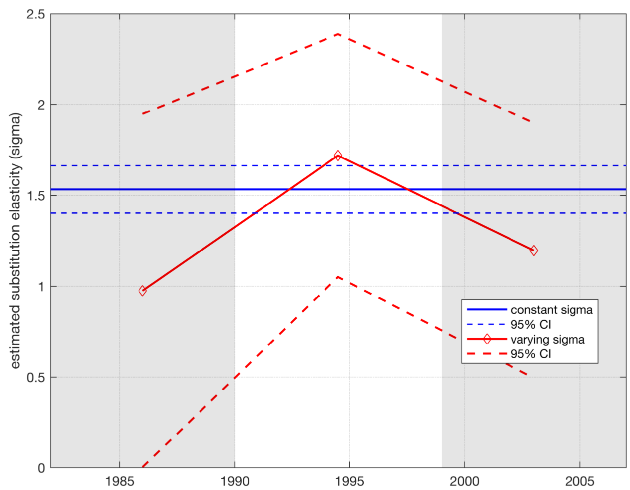

B.2 Additional details on capital-labor elasticity η

B.2.1 Markups and profits

Here I report on the relationship between labor shares and market power within 60 KLEMS industries over the period 1987–2015. As in the main text, I rely on the Lerner indices of market power (marginal profit shares) provided by Hall (2018). Hall estimates both industry-specific constants ($LernerLevel_i$) as well as linear time trends by industry ($LernerTrend_i$). I use labor shares in value added, and divide Lerner indices by the ratio of industry value added to industry sales. The latter transformation is conceptually appropriate, since labor shares are also measured as a fraction of value added, but the results are comparable without applying this transformation.

Column (1) of Table 13, as a basic sanity check, confirms that on average (pooled over all years), labor shares are lower in high market power industries. More relevant for this paper, Columns (2)-(4) report on the evolution of labor shares across industries. Column (2) says that on average, industry-level labor shares declined by 2.5 percentage points by decade (this number is very similar to the 2.3 p.p. decline per decade that I computed for goods based on the I-O Tables). Column (3) and (4) say that—across industries— there is a weak negative relation between labor share trends and the level of market power as measured by the Lerner indices, while increases in market power are not associated with falling labor shares (neither does the coefficient have the expected negative sign). I conclude that the findings reported in the main text, that falling labor shares are difficult to rationalize by potentially increasing market power, continue to hold within the KLEMS dataset, at the industry level.

Table 13: Labor shares and market power in KLEMS industries

	(1)	(2)	(3)	(4)
$LernerLevel_i$	-0.087**			
	(0.034)			
t		-0.025***	-0.022***	-0.026***
		(0.005)	(0.005)	(0.004)
$t \times LernerLevel_i$			-0.008*	
			(0.004)	
$t \times LernerTrend_i$				0.096
				(0.081)
Constant	0.685***	0.687***	0.687***	0.687***
	(0.045)	(0.006)	(0.006)	(0.006)
Industry i fixed effects	NO	YES	YES	YES

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Source: KLEMS, Hall (2018). The dependent variable is always the industry-level labor share θ^L_{it}, where i refers to one of 60 KLEMS industries. 29 time periods (1987–2015); $1,740$ observations. The unit of t is a decade. Standard errors in parentheses (clustered at industry level). All regressions weighted by industry value added.

B.2.2 Good-factor-specific technical progress

$$
\begin{aligned}
(\hat{\eta}-1) &= \frac{\operatorname{Cov}\left[\bar{\kappa}_i \hat{r}^E, (\eta-1)\bar{\kappa}_i \hat{r}^E - (\eta-1)\bar{\kappa}_i \hat{a}_i^E\right]}{\mathbb{V}\left[\bar{\kappa}_i \hat{r}^E\right]} \\
&= (\eta-1)\left(1 - \frac{\operatorname{Cov}\left[\bar{\kappa}_i \hat{r}^E, \bar{\kappa}_i \hat{a}_i^E\right]}{\mathbb{V}\left[\bar{\kappa}_i \hat{r}^E\right]}\right) \\
&= (\eta-1)\left(1 - \operatorname{Corr}\left[\bar{\kappa}_i \hat{r}^E, \bar{\kappa}_i \hat{a}_i^E\right] \frac{\sigma\left[\bar{\kappa}_i \hat{a}_i^E\right]}{\sigma\left[\bar{\kappa}_i \hat{r}^E\right]}\right) \\
&= (\eta-1)\left(1 + \operatorname{Corr}\left[\bar{\kappa}_i, \bar{\kappa}_i \hat{a}_i^E\right] \frac{\sigma\left[\bar{\kappa}_i \hat{a}_i^E\right]}{\sigma\left[\bar{\kappa}_i\right] \left|\hat{r}^E\right|}\right),
\end{aligned}
\tag{54}
$$

where the last line follows since $\hat{r}^E < 0$. $\hat{\eta}$ is an unbiased estimated of η, if $\bar{\kappa}_i$ is independent of $\hat{a}_i^E$, since then

$$
\operatorname{Cov}\left[\bar{\kappa}_i, \bar{\kappa}_i \hat{a}_i^E\right] = \mathbb{E}\left[\bar{\kappa}_i^2 \hat{a}_i^E\right] - \mathbb{E}\left[\bar{\kappa}_i\right] \mathbb{E}\left[\bar{\kappa}_i \hat{a}_i^E\right] = \mathbb{E}\left[\bar{\kappa}_i^2\right] \mathbb{E}\left[\hat{a}_i^E\right] - \mathbb{E}\left[\bar{\kappa}_i\right]^2 \mathbb{E}\left[\hat{a}_i^E\right] = 0, \tag{55}
$$

using independence of $\bar{\kappa}_i$ and $\hat{a}_i^E$, and moreover that $\mathbb{E}\left[\hat{a}_i^E\right] = 0$ (by definition, good-factor-specific technical progress is mean-zero).

B.2.3 Heterogeneous substitution elasticities

Suppose the true model is as in (35), but with η replaced by η_i. Then

$$(\hat{\eta}-1) = \frac{\text{Cov}\left[\bar{\kappa}_i(\hat{r}^E-\hat{r}^S),(\eta_i-1)\bar{\kappa}_i(\hat{r}^E-\hat{r}^S)\right]}{\mathbb{V}\left[\bar{\kappa}_i(\hat{r}^E-\hat{r}^S)\right]} = \frac{\text{Cov}\left[\bar{\kappa}_i,(\eta_i-1)\bar{\kappa}_i\right]}{\mathbb{V}\left[\bar{\kappa}_i\right]}. \tag{56}$$

The numerator in (56) can be re-written as

$$\begin{aligned}\text{Cov}\left[\bar{\kappa}_i,(\eta_i-1)\bar{\kappa}_i\right] &= \mathbb{E}\left[\bar{\kappa}_i^2(\eta_i-1)\right]-\mathbb{E}\left[\bar{\kappa}_i(\eta_i-1)\right]\mathbb{E}\left[\bar{\kappa}_i\right]\\ &= \text{Cov}\left[\bar{\kappa}_i^2,(\eta_i-1)\right]+\mathbb{E}\left[\bar{\kappa}_i^2\right]\mathbb{E}\left[(\eta_i-1)\right]-\left(\text{Cov}\left[\bar{\kappa}_i,(\eta_i-1)\right]+\mathbb{E}\left[\bar{\kappa}_i\right]\mathbb{E}\left[(\eta_i-1)\right]\right)\mathbb{E}\left[\bar{\kappa}_i\right]\\ &= \text{Cov}\left[\bar{\kappa}_i^2,(\eta_i-1)\right]-\text{Cov}\left[\bar{\kappa}_i,(\eta_i-1)\right]\mathbb{E}\left[\bar{\kappa}_i\right]+(\bar{\eta}-1)\left(\mathbb{E}\left[\bar{\kappa}_i^2\right]-\mathbb{E}\left[\bar{\kappa}_i\right]^2\right)\\ &= \left(\text{Cov}\left[\bar{\kappa}_i^2,(\eta_i-1)\right]-\text{Cov}\left[\bar{\kappa}_i,(\eta_i-1)\right]\mathbb{E}\left[\bar{\kappa}_i\right]\right)+(\bar{\eta}-1)\mathbb{V}\left[\bar{\kappa}_i\right].\end{aligned} \tag{57}$$

Plugging (57) into (56) yields

$$(\hat{\eta}-1) = (\bar{\eta}-1)+\frac{\text{Cov}\left[\bar{\kappa}_i^2,(\eta_i-1)\right]-\text{Cov}\left[\bar{\kappa}_i,(\eta_i-1)\right]\mathbb{E}\left[\bar{\kappa}_i\right]}{\mathbb{V}\left[\bar{\kappa}_i\right]}. \tag{58}$$

B.2.4 Relation to a framework with explicit input-output linkages

In this section, I show how the overall elasticity of substitution between capital and labor in final goods production, η, relates to the one for value added within an industry. Recall that η represents capital-labor substitution along the full value chain. To derive the relation between these two elasticities, I explicitly specify the input-output structure. Assume that good $i \in I = \{1, 2, ..., n\}$ is produced in industry i combining a value added bundle v_i with intermediate inputs x_i according to

$$y_i = \left((\alpha_i^y)^{\frac{1}{\eta_y}} v_i^{\frac{\eta_y-1}{\eta_y}} + (1-\alpha_i^y)^{\frac{1}{\eta_y}} x_i^{\frac{\eta_y-1}{\eta_y}}\right)^{\frac{\eta_y}{\eta_y-1}}, \text{ where } \alpha_i^y \in (0,1], \eta_y > 0 \tag{59}$$

$$v_i = \left((\alpha_i^v)^{\frac{1}{\eta_v}} k_i^{\frac{\eta_v-1}{\eta_v}} + (1-\alpha_i^v)^{\frac{1}{\eta_v}} l_i^{\frac{\eta_v-1}{\eta_v}}\right)^{\frac{\eta_v}{\eta_v-1}}, \text{ where } \alpha_i^v \in [0,1], \eta_v > 0 \tag{60}$$

$$x_i = \left(\sum_{j\in I}(\alpha_{ij}^x)^{\frac{1}{\eta_x}} q_{ij}^{\frac{\eta_x-1}{\eta_x}}\right)^{\frac{\eta_x}{\eta_x-1}}, \text{ where } \alpha_{ij}^x \geq 0, \sum_{j\in I}\alpha_{ij}^x = 1, \eta_x > 0. \tag{61}$$

In this framework, the three primitive elasticities are (i) the capital-labor elasticity within an industry's value added bundle (η_v), (ii) the "outsourcing" elasticity η_y that regulates how substitutable an industry's value added bundle is with intermediate inputs from other industries, (iii) and the intermediate input elasticity η_x that regulates how substitutable different intermediate inputs are. The resulting overall elasticity of substitution at the final good level, $\eta = (\eta_i)_{i\in I} \in \mathbb{R}^n$, is again defined as $\eta_i \equiv \frac{\partial \ln L_i/K_i}{\partial \ln R/W}$, where L_i and K_i refer to the inputs of labor and capital that are used in producing the final good y_i (not just in industry i, but in any industry). Note that in this framework, η will not be constant in general (unlike in the baseline model used in this paper), neither over time nor across goods.

Let p_i^y denote the price of good i, p_i^v the price of the value added bundle v_i, and p_i^x the price of the intermediate input bundle x_i. Throughout, I assume cost minimization, so these prices refer to optimal bundles. Let $\beta_i = \frac{p_i^v v_i}{p_i^y y_i}$ denote the value added cost share in production of good i, and $\Gamma_{ij} = \frac{p_j^y q_{i,j}}{p_i^x x_i}$ the good j cost share in the intermediate input bundle used for production of good i. Finally, let $\theta_i^{L,v} = \frac{Wl_i}{Wl_i+Rk_i} = \frac{Wl_i}{p_i^v v_i}$ denote the share of labor in value added of industry i. Then, the overall labor share of good i, θ_i^L, can be found by solving the following linear system:

$$\theta_i^L = \beta_i \theta_i^{L,v} + (1-\beta_i)\sum_{j\in I}\Gamma_{ij}\theta_j^L, \quad i = 1, ..., n. \tag{62}$$

Define the n-by-n matrix $\Gamma = [\Gamma_{ij}]_{ij}$, let $D(z)$ denote a matrix that has the entries of the vector z on the diagonal and zeroes off the main diagonal, and let I_n denote the identity matrix of size n. For all other objects defined earlier, z denotes the vector $(z_i)_{i\in I}$. Then the vector of overall final good labor shares can be solved for as

$$\theta^L = \left[I_n - D(\vec{1}-\beta)\Gamma\right]^{-1} D(\beta)\theta^{L,v} \tag{63}$$

Note that the matrix inverse is the well known Leontief inverse, which takes into account the infinite regress of industry A purchasing from industry B, which purchases from C, which purchases from A... Denoting by $\theta_i^{L,x}$ the labor share embodied in the intermediate input bundle x_i, we have that $\theta^{L,x} = \Gamma\theta^L$.

The following proposition relates η to the three underlying primitive elasticities η_v, η_y, η_x.

Proposition 2. *The overall capital-labor elasticity of substitution in production of good i,* $\eta \in \mathbb{R}^n$*, is a convex combination of* η_y*,* η_v *and* η_x*. In particular:*

$$\eta = M^{-1}\left(\eta_y a + \eta_v b + \eta_x c\right), \tag{64}$$

where

$$\begin{aligned}
M &= \left[I_n - D(\vec{1}-\beta)\Gamma\right] D(\theta^L)D(\vec{1}-\theta^L) \in \mathbb{R}^{n\times n}\\
a &= D(\beta)D(\vec{1}-\beta)D(\theta^{L,v}-\theta^{L,x})(\theta^{L,v}-\theta^{L,x}) \in \mathbb{R}^n\\
b &= D(\beta)D(\vec{1}-\theta^{L,v})\theta^{L,v} \in \mathbb{R}^n\\
c &= D(\vec{1}-\beta)\mathbb{V}^x(\theta^L) \in \mathbb{R}^n, \ \textit{where } \mathbb{V}^x(\theta^L) = \Gamma D(\theta^L)\theta^L - D(\theta^{L,x})\theta^{L,x}
\end{aligned}$$

Proof. First, note that the matrix version of (62) is

$$\theta^L = D(\beta)\theta^{L,v} + D(\vec{1}-\beta)\theta^{L,x} = D(\beta)\theta^{L,v} + D(\vec{1}-\beta)\Gamma\theta^L, \tag{65}$$

from which (63) follows immediately. Defining $\eta_i \equiv \frac{d\ln L_i/K_i}{d\ln R/W}$ as usual, we have that $\frac{d\theta_i^L}{d\ln R/W} = (1-\theta_i^L)\theta_i^L(\eta_i - 1)$ as before. Similarly, $\frac{d\theta_i^{L,v}}{d\ln R/W} = (1-\theta_i^{L,v})\theta_i^{L,v}(\eta_v - 1)$. Totally differentiate (62) with respect to $d\ln R/W$:

$$\begin{aligned}\frac{d\theta_i^L}{d\ln R/W} &= \frac{d\beta_i}{d\ln R/W}(\theta_i^{L,v} - \sum_{j\in I}\Gamma_{ij}\theta_j^L) + \beta_i\frac{d\theta_i^{L,v}}{d\ln R/W}\\ &\quad + (1-\beta_i)\sum_{j\in I}\frac{d\Gamma_{ij}}{d\ln R/W}\theta_j^L + (1-\beta_i)\sum_{j\in I}\Gamma_{ij}\frac{d\theta_j^L}{d\ln R/W}.\end{aligned} \tag{66}$$

To get expressions for $\frac{d\beta_i}{d\ln R/W}$ and $\frac{d\Gamma_{ij}}{d\ln R/W}$, first note that by Shephard's Lemma $\frac{d\ln p_i^v}{d\ln R/W} = (1-\theta_i^{L,v})$, $\frac{d\ln p_i^x}{d\ln R/W} = (1-\theta_i^{L,x})$ and $\frac{d\ln p_i^y}{d\ln R/W} = (1-\theta_i^L)$. Then,

$$\frac{d\ln\beta_i}{d\ln R/W} = (\eta_y - 1)\left(\frac{d\ln p_i^y}{d\ln R/W} - \frac{d\ln p_i^v}{d\ln R/W}\right) = (\eta_y - 1)\left(\theta_i^{L,v} - \theta_i^L\right), \tag{67}$$

$$\frac{d\ln\Gamma_{ij}}{d\ln R/W} = (\eta_x - 1)\left(\frac{d\ln p_i^x}{d\ln R/W} - \frac{d\ln p_j^y}{d\ln R/W}\right) = (\eta_x - 1)\left(\theta_j^L - \theta_i^{L,x}\right). \tag{68}$$

Substituting in for all derivatives in (66):

$$\begin{aligned}(1-\theta_i^L)\theta_i^L(\eta_i - 1) &= (\eta_y - 1)\beta_i\left(\theta_i^{L,v} - \theta_i^L\right)(\theta_i^{L,v} - \sum_{j\in I}\Gamma_{ij}\theta_j^L) + \beta_i(1-\theta_i^{L,v})\theta_i^{L,v}(\eta_v - 1)\\ &\quad + (1-\beta_i)\sum_{j\in I}(\eta_x - 1)\Gamma_{ij}\left(\theta_j^L - \theta_i^{L,x}\right)\theta_j^L + (1-\beta_i)\sum_{j\in I}\Gamma_{ij}(1-\theta_j^L)\theta_j^L(\eta_j - 1).\end{aligned} \tag{69}$$

(69) needs to be re-written in matrix form. Again, we have a linear system of n equations in n unknowns, namely the vector η:

$$\begin{aligned}D(\vec{1}-\theta^L)D(\theta^L)(\eta - 1) &= (\eta_y - 1)D(\beta)D(\theta^{L,v} - \theta^L)(\theta^{L,v} - \theta^{L,x}) + (\eta_v - 1)D(\beta)D(\vec{1}-\theta^{L,v})\theta^{L,v}\\ &\quad + (\eta_x - 1)D(\vec{1}-\beta)\left(\Gamma D(\theta^L)\theta^L - D(\theta^{L,x})\theta^{L,x}\right) + D(\vec{1}-\beta)\Gamma D(\vec{1}-\theta^L)D(\theta^L)(\eta - 1),\end{aligned} \tag{70}$$

which results in (64) when bringing the last term in (70) to the LHS and multiplying both sides by M^{-1} (from the left). The simplified expression for a in the proposition follows since

$$\begin{aligned}\beta_i\left(\theta_i^{L,v} - \theta_i^L\right)(\theta_i^{L,v} - \theta_i^{L,x}) &= \beta_i\left(\theta_i^{L,v} - \beta_i\theta_i^{L,v} - (1-\beta_i)\theta_i^{L,x}\right)(\theta_i^{L,v} - \theta_i^{L,x}) &(71)\\ &= \beta_i(1-\beta_i)(\theta_i^{L,v} - \theta_i^{L,x})^2. &(72)\end{aligned}$$

To proof the last claim, that η is indeed a convex combination of η_y, η_v and η_x, write out the

ith element of $a+b+c$, using repeatedly that $\theta_i^L = \beta_i \theta_i^{L,v} + (1-\beta_i)\theta_i^{L,x}$:

$$
\begin{aligned}
a_i + b_i + c_i &= \beta_i(1-\beta_i)((\theta_i^{L,v})^2 - 2\theta_i^{L,v}\theta_i^{L,x} + (\theta_i^{L,x})^2) + \beta_i\theta_i^{L,v} - \beta_i(\theta_i^{L,v})^2 \\
&\quad + (1-\beta_i)\sum_j \Gamma_{ij}(\theta_j^L)^2 - (1-\beta_i)(\theta_i^{L,x})^2 \qquad (73) \\
&= -\beta_i^2(\theta_i^{L,v})^2 - 2\beta_i(1-\beta_i)\theta_i^{L,v}\theta_i^{L,x} - (1-\beta_i)^2(\theta_i^{L,x})^2 \\
&\quad + \beta_i\theta_i^{L,v} + (1-\beta_i)\sum_j \Gamma_{ij}\theta_j^L(\theta_j^L + 1 - 1) \qquad (74) \\
&= -\left(\beta_i\theta_i^{L,v} + (1-\beta_i)\theta_i^{L,x}\right)^2 + \beta_i\theta_i^{L,v} + (1-\beta_i)\theta_i^{L,x} - (1-\beta_i)\sum_j \Gamma_{ij}\theta_j^L(1-\theta_j^L) \qquad (75) \\
&= -\left(\theta_i^L\right)^2 + \theta_i^L - (1-\beta_i)\sum_j \Gamma_{ij}\theta_j^L(1-\theta_j^L) \qquad (76) \\
&= \theta_i^L(1-\theta_i^L) - (1-\beta_i)\sum_j \Gamma_{ij}\theta_j^L(1-\theta_j^L), \qquad (77)
\end{aligned}
$$

or in vector notation

$$
a+b+c = \left(I_n - D(\vec{1}-\beta)\Gamma\right) D(\theta^L)(\vec{1}-\theta^L). \qquad (78)
$$

Hence

$$
\begin{aligned}
M^{-1}(a+b+c) &= D(\vec{1}-\theta^L)^{-1}D(\theta^L)^{-1}\left[I_n - D(\vec{1}-\beta)\Gamma\right]^{-1}\left(I_n - D(\vec{1}-\beta)\Gamma\right) D(\theta^L)(\vec{1}-\theta^L) \\
&= \vec{1}. \qquad (79)
\end{aligned}
$$

Moreover, a, b and c are all non-negative. □

Note that η is truly a convex combination of η_y, η_v, and η_x; i.e., $M^{-1}(a+b+c) = \vec{1}$ is an n-vector of ones, and the vectors a, b, and c are non-negative. Furthermore, note that besides the elasticities, all objects in this proposition are readily observable, respectively can easily be calculated, in an input-output table. The weight on η_y, for good i, is increasing in the difference between the labor share of the value added bundle v_i and of the intermediate input bundle x_i. Clearly, the value of η_y only matters if the labor shares of x_i and v_i are different. The weight on η_v is proportional to the cost share of the value added bundle, β_i. Finally, the weight on η_x is proportional to the weight on the intermediate input bundle times the variance of good-level labor shares weighted by the (equilibrium) intermediate input mix of good i. For η_x to affect η materially, it must be the case that intermediate inputs are quantitatively important and that labor shares of different intermediate inputs are sufficiently heterogeneous. Two polar cases are instructive: Without intermediate inputs, $\theta^L = \theta^{L,v}$ and $\beta = \vec{1}$, thus $a = c = \vec{0}$ and $\eta = \eta_v$. Of course, without intermediate inputs, the overall elasticity of substitution in production of good i is equal to the one for value added. To consider the other extreme, assume that industries use either only capital or only labor as direct inputs, i.e., $\theta_i^{L,v} \in \{0,1\}$. In that case, $b = \vec{0}$, and η is a convex combination of the outsourcing η_y as well as the intermediate input elasticity, η_x. Clearly, in that case the industry-level, value added, capital-labor elasticity η_v is irrelevant (it is not even well defined).

Table 14: Good-level (η) vs. industry-level (η_v) capital-labor elasticity

Dependent Variable:	$\log \theta_{it}^L$	$\log \left(\frac{\theta_{it}^L}{1-\theta_{it}^L} \right)$
	(1)	(2)
$(\eta - 1)$	0.427***	0.364**
	(0.133)	(0.158)
$(\eta_v - 1)$	0.278	0.168
	(0.256)	(0.335)

*** p<0.01, ** p<0.05, * p<0.1

All cells refer to IV estimates of $(\eta-1)$, resp. (η_v-1), weighted by final demand shares (η), resp. industry value added (η_v). Time and good fixed effects are used in all specifications; standard errors in parentheses (clustered at the good, respectively industry level). Columns (1) refers to equation (31), column (2) to equation (28). Equipment intensities for the manufacturing sector are taken from the NBER-CES manufacturing database; for non-manufacturing they are based on BEA-FAT.

It is possible, even simpler, to estimate η_v solely based on industry value added data, without taking into account the input-output structure. Table 14 compares the estimated industry-level elasticity η_v to the previously estimated overall good-level elasticity η. The results suggest that the overall elasticity η is indeed larger than the one for value added within an industry η_v. Such a relationship is expected if intermediate inputs are highly substitutable. As a case in point, Peter and Ruane (2018) estimate a plant-level elasticity of substitution between intermediate inputs of 4.3, using data on Indian manufacturing industries.[71] Their estimate refers to a 6-8 year horizon, comparable with this paper.

Irrespective of the relation between these two objects, if demand is specified in terms of final good expenditures as in this paper, η is the relevant object for questions such as the evolution of the aggregate labor share. In this case, estimating η is a more direct route, compared to estimating all three underlying elasticities.

[71] In the notation of this section, their estimate refers to the parameter η_x.

C Derivations and proofs

C.1 Proof of proposition 1

Proof. I drop time indices in this proof for clarity. Let the wage rate, in efficiency units, be the numeraire: $\hat{w} \equiv d\ln\frac{W}{A^L} = 0$. The market clearing condition for labor can be written as

$$W\bar{L} = E\left(\sum_{i\in I}\omega_i\theta_i^L\right), \tag{80}$$

total labor income equals total expenditure on labor. Take the logarithm and totally differentiate (80):

$$d\ln(W\bar{L}) = d\ln E + d\ln\left(\sum_{i\in I}\omega_i\theta_i^L\right) \tag{81}$$

$$\Rightarrow g_L = d\ln E + \frac{\sum_{i\in I}\omega_i\theta_i^L(d\ln\omega_i + d\ln\theta_i^L)}{\bar{\theta}^L}. \tag{82}$$

Observe that

$$d\ln E = d\ln(W\bar{L} + R\bar{K}) = \bar{\theta}^L g_L + (1-\bar{\theta}^L)(g_K + \hat{r}). \tag{83}$$

Furthermore, $d\ln p_i = (1-\theta_i^L)\hat{r} - g_A$ by Shephard's Lemma. Hence,

$$d\ln P \equiv \sum_{i\in I}\omega_i d\ln p_i = (1-\bar{\theta}^L)\hat{r} - g_A. \tag{84}$$

Thus, substituting in for $d\ln\frac{p_i}{P}$ and $d\ln\frac{E}{P}$ in (4),

$$d\ln\omega_i = (1-\sigma)(\bar{\theta}^L - \theta_i^L)\hat{r} + (\gamma_i - 1)g. \tag{85}$$

Substituting for $d\ln E$, $d\ln\omega_i$ and $d\ln\theta_i^L$ in (82):

$$\begin{aligned} g_L =& \bar{\theta}^L g_L + (1-\bar{\theta}^L)(g_K+\hat{r}) + \\ & \frac{\sum_{i\in I}\omega_i\theta_i^L\left((1-\sigma)(\bar{\theta}^L - \theta_i^L)\hat{r} + (\gamma_i-1)g + (\eta-1)(1-\theta_i^L)\hat{r}\right)}{\bar{\theta}^L}. \end{aligned} \tag{86}$$

Rearranging this equation, we can write it as

$$g_L - g_K = \tilde{\eta}\hat{r} + g\frac{Cov(\theta_i^L,\gamma_i)}{\bar{\theta}^L(1-\bar{\theta}^L)}, \tag{87}$$

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

$$\tilde{\eta} = \frac{\sigma \mathbb{V}[\theta_i^L] + \eta \mathbb{E}[\theta_i^L(1-\theta_i^L)]}{\bar{\theta}^L(1-\bar{\theta}^L)}. \tag{88}$$

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 the capital-labor mix is equal in all sectors, 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 of θ_i^L across sectors), then $\tilde{\eta} = \sigma$. The change in the aggregate labor share is given by $d\ln\bar{\theta}^L = g_L - d\ln E = (1-\bar{\theta}^L)(g_L - g_K - \hat{r})$. Substituting for $\hat{r}$ from (87) yields

$$d\bar{\theta}^L = \bar{\theta}^L d\ln\bar{\theta}^L = \bar{\theta}^L(1-\bar{\theta}^L)\left(g_L - g_K - \frac{1}{\tilde{\eta}}\left(g_L - g_k - g\frac{Cov(\theta_i^L,\gamma_i)}{\bar{\theta}^L(1-\bar{\theta}^L)}\right)\right), \tag{89}$$

which proves (9). □

C.2 Aggregation with log-linearity

Let $\theta_h^L = c_0 + c_1 \log(E_h)$ and $y_h = \log(E_h) \sim N(\mu - \frac{\sigma^2}{2}, \sigma^2)$. Then

$$\bar{\theta}^L \equiv \frac{\mathbb{E}[E_h\theta_h^L]}{\mathbb{E}[E_h]} = \frac{\mathbb{E}[E_h(c_0 + c_1\log(E_h))]}{\mathbb{E}[E_h]} = c_0 + c_1\frac{\mathbb{E}[E_h\log(E_h)]}{\mathbb{E}[E_h]} = c_0 + c_1\left(\mu + \frac{\sigma^2}{2}\right), \tag{90}$$

where the last equality follows since $\mathbb{E}[E_h] = \exp(\mu)$ and

$$\mathbb{E}[E_h\log(E_h)] = \mathbb{E}[\exp(y_h)y_h] \tag{91}$$

$$= \int \exp(y) y \frac{1}{\sqrt{2\sigma^2\pi}} \exp\left(-\frac{\left(y-\mu+\sigma^2/2\right)^2}{2\sigma^2}\right) dy \tag{92}$$

$$= \int y \frac{1}{\sqrt{2\sigma^2\pi}} \exp\left(y - \frac{y^2 - 2\mu y + y\sigma^2 - \mu\sigma^2 + \mu^2 + \sigma^4/4}{2\sigma^2}\right) dy \tag{93}$$

$$= \int y \frac{1}{\sqrt{2\sigma^2\pi}} \exp\left(-\frac{y^2 - 2y(\mu+\sigma^2/2) + (\mu+\sigma^2/2)^2 - 2\mu\sigma^2}{2\sigma^2}\right) dy \tag{94}$$

$$= \exp(\mu) \int y \frac{1}{\sqrt{2\sigma^2\pi}} \exp\left(-\frac{\left(y-(\mu+\sigma^2/2)\right)^2}{2\sigma^2}\right) dy \tag{95}$$

$$= \exp(\mu)(\mu + \sigma^2/2). \tag{96}$$

D Model with non-homothetic CES preferences

D.1 Preferences and demand system

In a static context, a consumer with nominal income E (= expenditure) maximizes an implicitly defined consumption aggregator U by choosing a consumption bundle $(c_i)_{i\in I}$:

$$\max \quad U \tag{97}$$

$$\text{subject to} \quad 1 = \sum_{i\in I} \Omega_i^{\frac{1}{\sigma}} c_i^{\frac{\sigma-1}{\sigma}} U^{\frac{\epsilon_i-\sigma}{\sigma}} \tag{98}$$

$$\text{and} \quad \sum_{i\in I} p_i c_i \leq E. \tag{99}$$

The following parametric restrictions are imposed: $\sigma > 0$ and $(\sigma-1)(\sigma-\epsilon_i) > 0 \ \forall \ i \in I$, the latter ensuring that $\frac{\partial U}{\partial c_i} > 0 \ \forall \ i$. Note that if $\epsilon_i = 1 \ \forall \ i$, the consumer has standard homothetic CES preferences.

Setting up the Lagrangian:

$$L = U - \lambda\left(1 - \sum_{i\in I} \Omega_i^{\frac{1}{\sigma}} c_i^{\frac{\sigma-1}{\sigma}} U^{\frac{\epsilon_i-\sigma}{\sigma}}\right) - \mu\left(\sum_{i\in I} p_i c_i - E\right). \tag{100}$$

The first-order condition with respect to c_i yields:

$$\lambda \frac{\sigma-1}{\sigma} \Omega_i^{\frac{1}{\sigma}} c_i^{\frac{-1}{\sigma}} U^{\frac{\epsilon_i-\sigma}{\sigma}} = \mu p_i. \tag{101}$$

Multiplying both sides by c_i, summing over all $i \in I$, and substituting in for the constraints,

$$\frac{\lambda}{\mu}\frac{\sigma-1}{\sigma} \underbrace{\sum_{i\in I} \Omega_i^{\frac{1}{\sigma}} c_i^{\frac{\sigma-1}{\sigma}} U^{\frac{\epsilon_i-\sigma}{\sigma}}}_{=1} = \underbrace{\sum_{i\in I} p_i c_i}_{=E}, \tag{102}$$

relates the Lagrange multipliers to E. Plugging this expression back into (101),

$$E \Omega_i^{\frac{1}{\sigma}} c_i^{\frac{-1}{\sigma}} U^{\frac{\epsilon_i-\sigma}{\sigma}} = p_i \Rightarrow c_i = \Omega_i U^{\epsilon_i-\sigma} E^{\sigma} p_i^{-\sigma}, \tag{103}$$

which expresses the optimal quantity of good i as a function of its price, total expenditure, and utility. Plugging this expression for c_i into the budget constraint yields the expenditure function $E = E((p_i)_{i\in I}, U)$:

$$E^{1-\sigma} = \sum_{i\in I} \Omega_i U^{\epsilon_i-\sigma} p_i^{1-\sigma}. \tag{104}$$

The map between utility and expenditure is smooth and strictly increasing as long as $(\sigma-1)(\sigma-\epsilon_i) > 0 \ \forall \ i$. It is, however, in general non-linear, unlike in the homothetic case.

Using the explicit definition of expenditure E as a function of utility U and prices in (104), and substituting in for E in (103) yields a closed-form solution of Hicksian demand $c_i((p_j)_{j\in I}, U)$:

$$c_i = \Omega_i p_i^{-\sigma} U^{\epsilon_i - \sigma} \left(\sum_{j\in I} \Omega_j U^{\epsilon_j-\sigma} p_j^{1-\sigma} \right)^{\frac{\sigma}{1-\sigma}}. \tag{105}$$

There is no closed-form solution for Marshallian demand $c_i((p_j)_{j\in I}, E)$. Given $((p_j)_{j\in I}, E)$, one can use (104) to solve for U and subsequently plug it in (103).

Denoting the expenditure share of good i by $\omega_i \equiv \frac{p_i c_i}{E} = \frac{p_i c_i}{\sum_{j\in I} p_j c_j}$, we have that

$$\omega_i = \frac{p_i c_i}{\sum_{j\in I} p_j c_j} = \frac{\Omega_i p_i^{1-\sigma} U^{\epsilon_i}}{\sum_{j\in I} \Omega_j p_j^{1-\sigma} U^{\epsilon_j}}, \tag{106}$$

or in terms of relative expenditure shares:

$$\frac{\omega_i}{\omega_j} = \frac{p_i c_i}{p_j c_j} = \frac{\Omega_i}{\Omega_j} \left(\frac{p_i}{p_j}\right)^{1-\sigma} U^{\epsilon_i - \epsilon_j}. \tag{107}$$

Substitution elasticities: The compensated substitution elasticity between any pair of goods is constant and equal to σ. Formally, when fixing U in (107) it immediately follows that

$$\frac{\partial \ln c_i/c_j}{\partial \ln p_j/p_i} = \sigma. \tag{108}$$

Income elasticities: From (107) it immediately follows that

$$\frac{\partial \ln c_i/c_j}{\partial \ln U} = \epsilon_i - \epsilon_j. \tag{109}$$

Relative demand elasticities, with respect to changes in utility U, are constant. However, U is an ordinal measure of utility. What are the implied income elasticities of Marshallian demand, i.e., $\frac{\partial \ln c_i}{\partial \ln E}|_{\text{fixed prices}}$, that we can estimate in the data? Differentiating the expenditure function (104) with respect to E:

$$(1-\sigma)E^{-\sigma} = \left(\sum_{i\in I} (\epsilon_i - \sigma) \Omega_i U^{\epsilon_i - \sigma - 1} p_i^{1-\sigma} \right) \frac{\partial U}{\partial E} \tag{110}$$

$$\Rightarrow (1-\sigma)E^{1-\sigma} = \left(\sum_{i\in I} (\epsilon_i - \sigma) \Omega_i U^{\epsilon_i - \sigma} p_i^{1-\sigma} \right) \frac{\partial \ln U}{\partial \ln E} \tag{111}$$

$$\Rightarrow (1-\sigma) = \left(\sum_{i\in I} (\epsilon_i - \sigma) \underbrace{\frac{\Omega_i U^{\epsilon_i-\sigma} p_i^{1-\sigma}}{\sum_{j\in I} \Omega_j U^{\epsilon_j - \sigma} p_j^{1-\sigma}}}_{=\omega_i} \right) \frac{\partial \ln U}{\partial \ln E} \tag{112}$$

$$\Rightarrow \frac{\partial \ln U}{\partial \ln E} = \frac{1-\sigma}{\bar{\epsilon}-\sigma}, \quad \text{where } \bar{\epsilon} \equiv \sum_{i\in I} \omega_i \epsilon_i. \tag{113}$$

Consequently, from (103)

$$\gamma_i \equiv \frac{\partial \ln c_i}{\partial \ln E} = \sigma + (\epsilon_i - \sigma)\frac{\partial \ln U}{\partial \ln E} = \sigma + (1-\sigma)\frac{\epsilon_i - \sigma}{\bar{\epsilon}-\sigma}. \tag{114}$$

Note that since $(\sigma-1)(\sigma-\epsilon_i) > 0 \ \forall \ i$, we have that $\frac{\epsilon_i-\sigma}{\bar{\epsilon}-\sigma} > 0$. Thus, non-homothetic CES preferences bound (Marshallian) income elasticities: if $\sigma < 1$, then $\gamma_i > \sigma$ for all i, i.e., income elasticities are bounded from below. Of course, from the budget constraint the expenditure-weighted average income elasticity has to equal one (formally $\bar{\gamma} \equiv \sum_{i\in I}\omega_i\gamma_i = 1$). Analoguesly, if $\sigma > 1$, then $\gamma_i < \sigma$ for all i. Thus, the closer σ is to one, the tighter the bound on income elasticities imposed by these preferences. In the limit, as $\sigma \to 1$, the implied demand system converges to the one implied by Cobb-Douglas preferences, not only with respect to the implied substitution elasticity, but also with respect to the implied income elasticities.

First-order approximation to Marshallian demand system: In section 3, I proposed a simple ad hoc demand system in equation (4). Here, I show how this demand system can be rationalized as first-order log-linear approximation to the one implied by non-homothetic CES preferences. Note that—both in section 3 as well as here—$d\ln P \equiv \sum_{i\in I}\omega_i d\ln p_i$ is the empirically-motivated definition of an expenditure-weighted price deflator, not the ideal price index implied by non-homothetic CES preferences. Start from (107), log-linearize, and subsequently sum over all $j \in I$, weighted by ω_j:

$$d\ln\omega_i - d\ln\omega_j = (1-\sigma)(d\ln p_i - d\ln p_j) + (\epsilon_i - \epsilon_j)d\ln U \tag{115}$$

$$d\ln\omega_i - \underbrace{\sum_{j\in I}\omega_j d\ln\omega_j}_{=0} = (1-\sigma)(d\ln p_i - \underbrace{\sum_{j\in I}\omega_j d\ln p_j}_{=d\ln P}) + (\epsilon_i - \underbrace{\sum_{j\in I}\omega_j\epsilon_j}_{=\bar{\epsilon}})d\ln U \tag{116}$$

$$d\ln\omega_i = (1-\sigma)d\ln\frac{p_i}{P} + (\epsilon_i - \bar{\epsilon})d\ln U. \tag{117}$$

Then, totally differentiate the expenditure function (104):

$$(1-\sigma)E^{1-\sigma}d\ln E = \sum_{i\in I}\Omega_i U^{\epsilon_i-\sigma}p_i^{1-\sigma}\left((1-\sigma)d\ln p_i + (\epsilon_i-\sigma)d\ln U\right) \tag{118}$$

$$\Rightarrow (1-\sigma)d\ln E = \sum_{i\in I}\underbrace{\frac{\Omega_i U^{\epsilon_i-\sigma}p_i^{1-\sigma}}{\sum_{j\in I}\Omega_j U^{\epsilon_j-\sigma}p_j^{1-\sigma}}}_{=\omega_i}\left((1-\sigma)d\ln p_i + (\epsilon_i-\sigma)d\ln U\right) \tag{119}$$

$$\Rightarrow d\ln E = \sum_{i\in I}\omega_i d\ln p_i + \frac{\sum_{i\in I}\omega_i\epsilon_i - \sigma}{1-\sigma}d\ln U \tag{120}$$

$$\Rightarrow d\ln U = \frac{1-\sigma}{\bar{\epsilon}-\sigma}d\ln\frac{E}{P}. \tag{121}$$

Combining (117) and (121), we confirm that

$$d\ln\omega_i = (1-\sigma)d\ln\frac{p_i}{P} + (\epsilon_i - \bar{\epsilon})\frac{1-\sigma}{\bar{\epsilon}-\sigma}d\ln\frac{E}{P} \tag{122}$$

$$\Rightarrow d\ln\omega_i = (1-\sigma)d\ln\frac{p_i}{P} + (\gamma_i - 1)d\ln\frac{E}{P}, \tag{123}$$

where γ_i is the Marshallian income elasticity as in (114).

D.2 Estimating income elasticities with non-homothetic CES preferences

The starting point is the expression for the expenditure share (107) in time t, for household h, and good i, relative to the reference good 0. Taking logs yields

$$\ln\left(\frac{\omega_{iht}}{\omega_{0ht}}\right) = \ln\left(\frac{\Omega_i}{\Omega_0}\right) + (1-\sigma)\ln\left(\frac{p_{it}}{p_{0t}}\right) + (\epsilon_i - \epsilon_0)\ln(U_{ht}). \tag{124}$$

The aim is to estimate the income elasticity parameters ϵ_i from cross-sectional variation in household expenditure. However, utility U_{ht} is not observed. The difficulty is that U_{ht} is an only implicitly defined (but strictly increasing) function of observed household expenditure E_{ht}, which depends also on the unknown demand parameters $((\Omega_i)_{i\in I}, \sigma)$ in addition to prices.

Proceeding in two steps, I first estimate

$$\ln\left(\frac{\omega_{iht}}{\omega_{0ht}}\right) = \tilde{\zeta}_{it} + \beta_{it}\ln(E_{ht}) + \Gamma'_{it}Z_{ht} + \xi_{iht} \tag{125}$$

separately for each i and t. That is, I simply replace U_{ht} by E_{ht}. Having estimated $\hat{\beta}_{it}$ for good i for every t, I define $\hat{\beta}_i$ as the time-average.

In the second step, I map the reduced form relative income elasticity parameters β_i into the structural parameters ϵ_i.This second step requires knowing the substitution elasticity σ, which will be solved for as a fixed point. Specifically, I employ a log-linear approximation of the inverse expenditure function (indirect utility function for fixed prices) around average expenditure $\bar{E}_t$ as in (113):

$$\ln(U_{ht}) = \ln(\bar{U}_t) + \frac{1-\sigma}{\bar{\epsilon}_t - \sigma}(\ln(E_{ht}) - \ln(\bar{E}_t)) + O((\ln(E_{ht}) - \ln(\bar{E}_t))^2), \tag{126}$$

where $\bar{\epsilon}_t = \sum_i \bar{\omega}_{it}\epsilon_i$. Thus, locally around $\bar{E}_t$ it holds that

$$\beta_{it} = \frac{\partial\ln\left(\frac{\omega_{iht}}{\omega_{0ht}}\right)}{\partial\ln(E_{ht})} = \frac{\partial\ln\left(\frac{\omega_{iht}}{\omega_{0ht}}\right)}{\partial\ln(U_{ht})}\frac{\partial\ln(U_{ht})}{\partial\ln(E_{ht})} = (\epsilon_i - \epsilon_0)\frac{1-\sigma}{\bar{\epsilon}_t - \sigma}. \tag{127}$$

Hence,

$$\epsilon_i = \epsilon_0 + \beta_{it}\frac{\bar{\epsilon}_t - \sigma}{1-\sigma}. \tag{128}$$

Plugging this into the definition of $\bar{\epsilon}_t$ yields

$$\bar{\epsilon}_t = \sum_i \bar{\omega}_{it}\epsilon_i = \sum_i \bar{\omega}_{it}(\epsilon_0 + \beta_i \frac{\bar{\epsilon}_t - \sigma}{1-\sigma}) = \epsilon_0 + \bar{\beta}_t \frac{\bar{\epsilon}_t - \sigma}{1-\sigma}, \tag{129}$$

where $\bar{\beta}_t = \sum_i \bar{\omega}_{it}\beta_{it}$ is the average income elasticity relative to good 0. This is a linear equation relating ϵ_0 to $\bar{\epsilon}_t$, which we can solve for

$$\epsilon_0 = \bar{\epsilon}_t \frac{1-\sigma-\bar{\beta}_t}{1-\sigma} + \frac{\sigma\bar{\beta}_t}{1-\sigma}. \tag{130}$$

Finally, plugging (130) into (128) we get an expression for the structural parameters ϵ_i as a function of $\bar{\epsilon}_t$, σ, and estimated reduced form parameters:

$$\epsilon_i = -\sigma\frac{\beta_i - \bar{\beta}_t}{1-\sigma} + \bar{\epsilon}_t\left(1 + \frac{\beta_i - \bar{\beta}_t}{1-\sigma}\right). \tag{131}$$

Given a value of σ, the ϵ_i are identified up to one degree of freedom, their level $\bar{\epsilon}_t$.

It turns out that given $(\beta_i)_{i\in I}$ and $\bar{\beta}_t$, there is a set $\Sigma \subset \mathbb{R}_{++}$ such that a solution consistent with the parameter restrictions exists if and only if $\sigma \in \Sigma$. A solution here means an $\bar{\epsilon}_t$ that gives rise to $(\epsilon)_{i\in I}$ that satisfy the restriction that $(\sigma-1)(\sigma-\epsilon_i) > 0\ \forall i \in I$. Furthermore, either $\Sigma = (0, \underline{\sigma}) \cup (\bar{\sigma}, +\infty)$ or $\Sigma = (\bar{\sigma}, +\infty)$ for some $\underline{\sigma} \in (0,1)$ and some $\bar{\sigma} > 1$.

To see this, note that at $\bar{\epsilon}_t = \sigma$, it holds that $\epsilon_i = \sigma$ for all i. Also note that the ϵ_i are linear in $\bar{\epsilon}_t$. Since $\bar{\epsilon}_t$ is an average over the ϵ_i, for some i it must hold that ϵ_i is increasing in $\bar{\epsilon}_t$. Since $(\sigma - \epsilon_i)$ must have the same sign for all i, for all i it must hold that ϵ_i increases in $\bar{\epsilon}_t$; i.e., $\left(1 + \frac{\beta_i - \bar{\beta}_t}{1-\sigma}\right) > 0\ \forall i$. If $\sigma < 1$, then this condition holds for all i if it holds for the smallest relative income elasticity β_{min}. In turn, this defines $\underline{\sigma} = 1 + (\beta_{min} - \bar{\beta}_t) < 1$. Else, if $\sigma > 1$, then the condition holds for all i if it holds for the largest relative income elasticity β_{max}. In turn, this defines $\bar{\sigma} = 1 + (\beta_{max} - \bar{\beta}_t) > 1$.

This, of course, is another way of stating the earlier observation that non-homothetic CES preferences do not allow for having both strong income effects and weak price effects.

D.3 Further approach

Given a guess for σ, it is convenient to normalize utility $U_\tau = 1$ in the base period ($\tau = 1982$). Then the demand level parameters Ω_i can be set to equal expenditure shares, i.e., $\Omega_i = \bar{\omega}_{i\tau}^{CEX}$.[72]

[72]Note that model expenditure shares are given by $\omega_{i\tau} = \frac{\Omega_i p_{i\tau}^{1-\sigma} U_\tau^{\epsilon_i}}{\sum_{j\in I}\Omega_j p_{j\tau}^{1-\sigma} U^{\epsilon_j}} = \frac{\Omega_i}{\sum_{j\in I}\Omega_j} = \bar{\omega}_{i\tau}^{CEX}$. The expenditure function indeed evaluates to one, $E_\tau = \left(\sum_{i\in I}\Omega_i U_\tau^{\epsilon_i - \sigma} p_i^{1-\sigma}\right)^{\frac{1}{1-\sigma}} = \left(\sum_{i\in I}\bar{\omega}_{i\tau}^{CEX}\right)^{\frac{1}{1-\sigma}} = 1$.

D.4 Substitution elasticity

Under non-homothetic CES demand, solving the model yields a time series of utility U_t. Rearranging the FOCs for demand:

$$\ln\left(\frac{\bar{\omega}_{it}^{CEX}}{\bar{\omega}_{0t}^{CEX}}\right) - \ln\left(\frac{\Omega_i}{\Omega_0}\right) - (\epsilon_i - \epsilon_0)\ln(U_t) = (1-\sigma)\ln\left(\frac{p_{it}}{p_{0t}}\right). \tag{132}$$

All terms (including the utility function parameters) on the LHS are known. I estimate σ by regressing relative prices on residual demand (pooling all goods and years).

As explained in Section D.1, these preferences impose a joint restriction on income and substitution elasticities. It turns out that this restriction is binding. Thus, the preference system is rejected. Specifically, the estimated income elasticities imply $\underline{\sigma} = -0.59$ and $\bar{\sigma} = 2.34$; i.e., $\sigma \geq 2.34$. When calibrating the model, the estimation routine is stuck at the corner $\bar{\sigma}$: when guessing $\sigma = 2.34$ and solving the model, regression (132) returns $\hat{\sigma} < \bar{\sigma}$. This is not surprising as the reduced form demand specification that is used in the main text returns $\hat{\sigma} = 1.51$. In general, if one were to estimate income and substitution elasticity parameters jointly, the estimated parameters would give rise to a combination of income and substitution elasticities that does not match the data (i.e., not even locally).

D.5 Results

The results of the model with non-homothetic CES utility are displayed in Figures 19. The only difference to the specification in the main text is that, since σ is constrained to be larger, substitution towards capital-intensive goods is stronger. Therefore, the labor share decline is steeper: over the full time period, the fall is overstated by 48.6%.

Figure 19: Non-homothetic CES preferences: Aggregate labor share

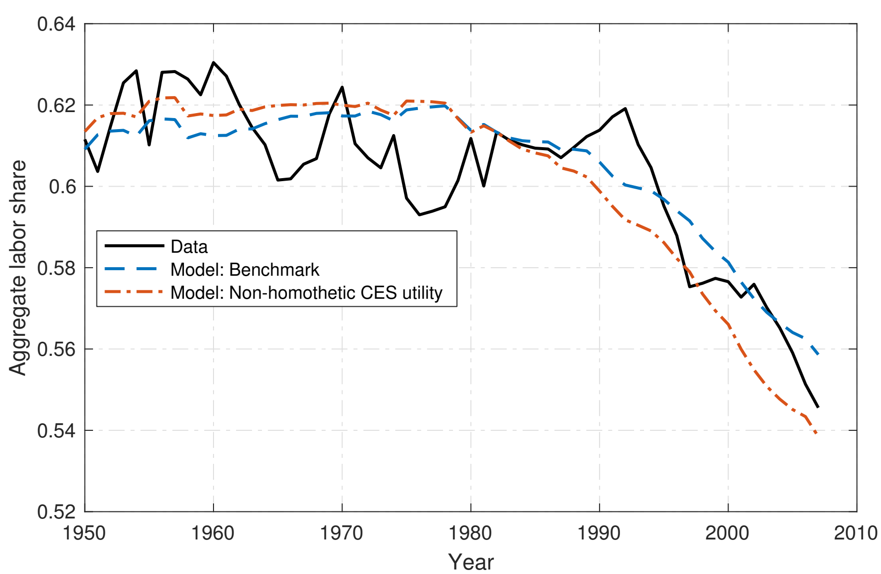

References

Acemoglu, Daron and Pascual Restrepo, "Robots and Jobs: Evidence from US Labor Markets," Working Paper 23285, National Bureau of Economic Research March 2017.

_ **and** _ , "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, June 2018, *108* (6), 1488–1542.

Aguiar, Mark and Mark Bils, "Has Consumption Inequality Mirrored Income Inequality?," *American Economic Review*, September 2015, *105* (9), 2725–56.

Antràs, Pol, "Is the U.S. Aggregate Production Function Cobb-Douglas? New Estimates of the Elasticity of Substitution," *The B.E. Journal of Macroeconomics*, April 2004, *4* (1), 1–36.

Attanasio, Orazio P. and Luigi Pistaferri, "Consumption Inequality," *Journal of Economic Perspectives*, May 2016, *30* (2), 3–28.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen, "The Fall of the Labor Share and the Rise of Superstar Firms," CEP Discussion Papers dp1482, Centre for Economic Performance, LSE May 2017.

Barkai, Simcha, "Declining Labor and Capital Shares," Technical Report, University of Chicago 2017.

Becker, Randy, Wayne Gray, and Jordan Marvakov, "NBER-CES Manufacturing Industry Database: Technical Notes," Technical Report, National Bureau of Economic Research 2016.

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

Borusyak, Kirill and Xavier Jaravel, "The Distributional Effects of Trade: Theory and Evidence from the United States," 2018 Meeting Papers 284, Society for Economic Dynamics 2018.

Buera, Francisco J. and Joseph P. Kaboski, "Can Traditional Theories of Structural Change Fit The Data?," *Journal of the European Economic Association*, 04-05 2009, *7* (2-3), 469–477.

Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas, "Rents, Technical Change, and Risk Premia Accounting for Secular Trends in Interest Rates, Returns on Capital, Earning Yields, and Factor Shares," *American Economic Review*, May 2017, *107* (5), 614–20.

Chirinko, Robert S., "σ: The long and short of it," *Journal of Macroeconomics*, 2008, *30* (2), 671 – 686. The CES Production Function in the Theory and Empirics of Economic Growth.

Comin, Diego A, Danial Lashkari, and Mart Mestieri, "Structural Change with Long-run Income and Price Effects," Working Paper 21595, National Bureau of Economic Research September 2015.

Cummins, Jason G. and Giovanni L. Violante, "Investment-Specific Technical Change in the US (1947-2000): Measurement and Macroeconomic Consequences," *Review of Economic Dynamics*, April 2002, *5* (2), 243–284.

DiCecio, Riccardo, "Sticky wages and sectoral labor comovement," *Journal of Economic Dynamics and Control*, March 2009, *33* (3), 538–553.

Eden, Maya and Paul Gaggl, "On the Welfare Implications of Automation," *Review of Economic Dynamics*, July 2018, *29*, 15–43.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu, "How Costly Are Markups?," Working Paper 24800, National Bureau of Economic Research July 2018.

Elsby, Michael, Bart Hobijn, and Aysegül Sahin, "The Decline of the U.S. Labor Share," Working Paper Series 2013-27, Federal Reserve Bank of San Francisco September 2013.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer, "The Next Generation of the Penn World Table," *American Economic Review*, October 2015, *105* (10), 3150–3182.

Fisher, Jonas D. M., "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks," *Journal of Political Economy*, June 2006, *114* (3), 413–451.

Garner, Thesai, George Janini, William Passero, Laura Paszkiewicz, and Mark Vendemia, "The Consumer Expenditure Survey: A Comparison with Personal Expenditures.," Technical Report, Bureau of Labor Statistics 2006.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," Working Paper 24408, National Bureau of Economic Research March 2018.

Gollin, Douglas, "Getting Income Shares Right," *Journal of Political Economy*, 2002, *110* (2), 458–474.

Gordon, Robert J, *The Measurement of Durable Goods Prices*, Chicago, IL: University of Chicago Press, 1990.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, "Long-Run Implications of Investment-Specific Technological Change," *The American Economic Review*, 1997, *87* (3), 342–362.

Hall, Robert E, "Using Empirical Marginal Cost to Measure Market Power in the US Economy," Working Paper 25251, National Bureau of Economic Research November 2018.

Herrendorf, Berthold, Richard Rogerson, and kos Valentinyi, "Two Perspectives on Preferences and Structural Transformation," *American Economic Review*, December 2013, *103* (7), 2752–89.

_ , _ , **and** _ , "Structural Change in Investment and Consumption: A Unified Approach," Working Paper 24568, National Bureau of Economic Research May 2018.

Horowitz, Karen J. and Mark A. Planting, *Concepts and Methods of the U.S. Input-Output Accounts*, U.S. Department of Commerce, Bureau of Economic Analysis, 2014.

Hulten, Charles R., "Growth Accounting When Technical Change is Embodied in Capital," *The American Economic Review*, 1992, *82* (4), 964–980.

Kaldor, Nicholas, "Capital Accumulation and Economic Growth," in D. C. Hague, ed., *The Theory of Capital: Proceedings of a Conference held by the International Economic Association*, London: Palgrave Macmillan UK, 1961, pp. 177–222.

Karabarbounis, Loukas and Brent Neiman, "The Global Decline of the Labor Share," *The Quarterly Journal of Economics*, 2014, *129* (1), 61–103.

_ **and** _ , "Accounting for Factorless Income," NBER Working Papers 24404, National Bureau of Economic Research, Inc March 2018.

Kehrig, Matthias and Nicolas Vincent, "Growing Productivity without Growing Wages: The Micro-Level Anatomy of the Aggregate Labor Share Decline," CESifo Working Paper Series 6454, CESifo Group Munich 2017.

Koh, Dongya, Ral Santaeullia-Llopis, and Yu Zheng, "Labor Share Decline and Intellectual Property Products Capital," Working Papers 927, Barcelona Graduate School of Economics September 2016.

León-Ledesma, Miguel A., Peter McAdam, and Alpo Willman, "Identifying the Elasticity of Substitution with Biased Technical Change," *American Economic Review*, September 2010, *100* (4), 1330–57.

Levinson, Arik and James O'Brien, "Environmental Engel Curves," NBER Working Papers 20914, National Bureau of Economic Research, Inc January 2015.

Loecker, Jan De and Jan Eeckhout, "The Rise of Market Power and the Macroeconomic Implications," NBER Working Papers 23687, National Bureau of Economic Research, Inc August 2017.

Oberfield, Ezra and Devesh Raval, "Micro Data and Macro Technology," NBER Working Papers 20452, National Bureau of Economic Research, Inc September 2014.

Peter, Alessandra and Cian Ruane, "The Aggregate Importance of Intermediate Input Substitutability," Working Paper January 2018.

Piketty, Thomas, *Capital in the Twenty-First Century*, Translated by Arthur Goldhammer. Cambridge, MA: Belknap, 2014.

Rognlie, Matthew, "Deciphering the Fall and Rise in the Net Capital Share: Accumulation or Scarcity?," *Brookings Papers on Economic Activity*, 2015, *46* (1 (Spring), 1–69.

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Nicholas Trachter, "Diverging Trends in National and Local Concentration," Working Paper 25066, National Bureau of Economic Research September 2018.

Traina, James, "Is Aggregate Market Power Increasing? Production Trends Using Financial Statements," Working Paper 2018.

Valentinyi, Akos and Berthold Herrendorf, "Measuring Factor Income Shares at the Sector Level," *Review of Economic Dynamics*, October 2008, *11* (4), 820–835.