Consumption Inequality in the Digital Age*

Kai Arvai†
European Central Bank

Katja Mann‡
Copenhagen Business School

12 March 2026

Abstract

This paper studies how digitalization affects consumption inequality. We assemble a novel dataset of digital technology in consumption, and establish that high-income households consume a larger share of digitally produced products than low-income households. We show that these products become relatively cheaper, which benefits rich households and contributes to consumption inequality. In a structural model with heterogeneous workers and non-homothetic preferences, we quantify the impact of digitalization on consumption inequality. The model demonstrates that a U-shaped income polarization translates into J-shaped welfare changes, as the price effect mutes income gains at the bottom while amplifying them at the top.

Keywords: Digitalization, Inequality, Consumption, Automation, Relative prices

JEL Classification: E21, E22, J31, O33, O41

*We would like to thank four anonymous referees, Melanie Arntz, Christian Bayer, Andrea Colciago, Gabriele Guaitoli, Yannick Kalantzis, Karin Kinnerud, Keith Kuester, Moritz Kuhn, Agnieszka Markiewicz, Morten Olsen, Klaus Prettner, Pontus Rendahl, Juhana Siljander and Christoffer Weissert for helpful comments and suggestions. We are grateful to Xavier Jaravel and Kirill Borusyak for sharing their CEX and BLS concordances and to Joachim Hubmer for sharing his production dataset. A big thanks also goes to Pedro Salgado for excellent research assistance. The views in this paper are the ones of the authors and do not represent the views of the European Central Bank or the Eurosystem.

†European Central Bank, E-mail address: kai.arvai@ecb.europa.eu

‡Copenhagen Business School, E-mail address: kma.eco@cbs.dk

1 Introduction

Digital technology transforms our economy, as it fundamentally changes the way we produce and consume. The increased usage of computers and other digital assets in production has resulted in wage and employment polarization in the Western world (Michaels et al., 2014; Akerman et al., 2015; Gaggl and Wright, 2017; Burstein et al., 2019). Over the last decades, the top and the bottom of the income distribution have gained in terms of income and employment shares, whereas jobs in the middle of the distribution have declined. However, progress in digital technology may not only have an *income effect*, but also a *price effect*: Sectors that rely more on digital assets can potentially produce at a lower cost and offer lower prices to their customers. As consumption patterns differ across the income distribution, so do price changes. Depending on the direction of this effect, price changes could either reinforce or counteract income changes in welfare calculations. This paper measures the price effect and shows how digitalization affects consumption inequality as well as overall welfare in the United States.

By combining household- and sector-level data, we establish that the digital share of the consumption basket monotonically increases along the income distribution. We show that the relative price decline of digital products and services, together with the difference in consumption expenditures, benefits richer households more and explains a substantial part of observed inequality in consumer prices. To assess the importance of the price effect for consumption and welfare, in particular compared to the income effect, we present a dynamic general equilibrium model with heterogeneous workers and non-homothetic preferences, which we calibrate to the US economy between 1960 and 2017. When taking income-specific price changes into account, the well-known U-shape of income changes translates into a J-shape for welfare: At the top, income gains from digitalization get amplified through favorable price dynamics, whereas they get muted at the bottom.

A priori it is unclear how price changes induced by digitalization impact consumption inequality. As the increased use of digital technology makes some consumption goods cheaper than others, it will benefit the income groups that consume relatively more of these goods. Depending on what those goods are, either rich or poor households could be the beneficiaries. Identifying the direction of the price effect is therefore first and foremost an empirical task, which requires a measure of the digital content of goods and services. Using data from the U.S. Bureau of Economic Analysis (BEA), we identify capital goods that relate to Information and Communication Technology (ICT) and estimate their rate of return and their share in value added. We document that the share of ICT capital in value-added has increased substantially between 1960 and 2017. Importantly, there is large heterogeneity in the usage of ICT capital across industries. The information industry and the computer manufacturing industry, but also some financial services, are ICT-intensive, whereas some other manufacturing industries and services are less digitalized.

We account for the digital content of intermediate products by leveraging the input-output structure of the production network. The resulting measure covers about 400 final goods and services. We link these to consumption categories in the Consumer Expenditure Survey (CEX). Using CEX data on household income, we construct the ICT share of consumption baskets along the income distribution. We find that the ICT share of consumption increases with income. Rich households

have a 20% larger digital share than poor households. We also study how the time-use for ICT-intensive activities is distributed across the income distribution, and find a positive correlation with income. Additionally, we discuss the implications of measuring ICT differently.

In the next step, we establish a link between digitalization and changes in consumer prices. We document that goods and services with a larger digital share have experienced lower price increases, confirming conventional wisdom about the impact of technological progress on prices. Since richer households consume a larger share of digital-intensive goods, the price of their consumption basket increases less. This aligns with existing empirical evidence on inflation inequality and establishes digitalization as a candidate explanation for observed differences in price changes across the income distribution.

In the second part of the paper, we present a structural model to assess the welfare effect of digitalization for different income groups. While the data provide a first estimate of the price effect, the model generates changes in both input prices and wages and allows for substitution in input factors, which are important for final output prices. Additionally, the model enables us to assess welfare effects, which cannot be measured empirically. The model features a two-sector economy with two types of capital, ICT and non-ICT. Sector 2 uses ICT capital more intensively. The economy is populated by two types of households, high-skill and low-skill, where low-skill households are heterogeneous and sort endogenously into routine and non-routine manual occupations. In this set-up, we follow Jaimovich et al. (2021).

Digitalization is modeled as an increase in the rate of transformation of output into ICT capital, which affects on the one hand the factor demands and on the other hand the relative price of the two goods. There are several other exogenous processes next to digitalization: sector-specific TFP growth, changes in the productivity of non-ICT capital, and changes in labor supply. We combine this setup with non-homothetic preferences as in Boppart (2014), implying that the effect of changing relative prices depends on the position in the income distribution.

We calibrate the model to the U.S. economy between 1960 and 2017. We use the method of simulated moments to match key moments in the data, in particular the increase in the skill premium, the decline in the relative price of ICT-intensive goods and differences in consumption between income groups. The model generates income polarization and heterogeneous price dynamics that are very close to the data.

The income effect is about as large for households at the top and bottom of the distribution, while being smaller in the middle. The price effect shows a different pattern. The richer the household, the higher the relative gain in purchasing power from digitalization. As a result, the overall welfare effects vary dramatically. The U-shaped polarization in income transforms into a J-shape in welfare as income gains at the bottom get muted and at the top get amplified by the price effect. The price effect is large especially for the very top and bottom of the income distribution, changing the relative welfare changes caused by income polarization by 30%. While in our set-up, differences in price changes thus lead to differences in welfare, Oberfield (2023) shows theoretically that this does not necessarily have to be the case. We discuss the differences between his model and ours in the context of our results.

To isolate the effect of digitalization on incomes, prices and welfare, we conduct a counterfac-

tual exercise in which we mute advances in digital technology. Income polarization disappears, whereas the price effect is essentially flat. This highlights the important role of digitalization for both channels. In a second counterfactual exercise, we show that in absence of non-homothetic preferences, the expenditure shares would be the same along the income distribution, and overall welfare would follow the evolution of income one-to-one. Non-homothetic preferences are therefore key for the overall welfare effects.

We make two contributions to the literature. First, we identify the technology content of consumption goods[1] and establish digital technology as a driver of relative price and consumption inequality. Several recent papers have documented that inflation rates are higher for low-income than for high-income households (Kaplan and Schulhofer-Wohl, 2017; Jaravel, 2019; Argente and Lee, 2021; Jaravel and Lashkari, 2024). Less is known about the drivers of this pattern. While Jaravel (2019) argues for a role of innovation, studies on the role of trade provide contradictory results. Fajgelbaum and Khandelwal (2016) find trade to lead to lower inflation rates for the poor, thus going in the opposite direction, whereas Borusyak and Jaravel (2021) find it to be flat along the income distribution. We provide evidence that differences in the digital content of the consumption basket can generate substantial inequality in relative prices. In studying the heterogeneous effect of digitalization on consumption, we also contribute to the literature documenting and explaining differences in consumption patterns between rich and poor households (Krueger and Perri, 2006; Attanasio and Pistaferri, 2014; Aguiar and Bils, 2015; Meyer and Sullivan, 2023). To our knowledge, we are the first to relate consumption inequality to technological trends, both theoretically and empirically. We also highlight the importance of distinguishing between nominal and real consumption inequality.

Second, we provide a comprehensive assessment of the welfare effects of digitalization and automation. While many papers have documented a U-shaped polarization pattern focusing on changes in employment and wages (Goos and Manning, 2007; Goos et al., 2009; Autor et al., 2008; Autor and Dorn, 2013; Michaels et al., 2014), or wealth (Prettner and Strulik, 2020; Moll et al., 2022), we establish the price effect as another crucial driver of welfare effects. To a relatively standard framework on the income side (see e.g. Eden and Gaggl, 2018; Jaimovich et al., 2021; Vom Lehn, 2020; Bárány and Siegel, 2021), we introduce multiple consumption goods and non-homothetic preferences to give rise to the price channel. We show that this channel is sizable if matched with empirical moments and that it redefines the winners and losers from automation.

Section 2 presents the empirical analysis. In Section 3 we introduce our model. Section 4 explains the calibration, followed by the simulation results in Section 5. Section 6 concludes.

2 Empirical Analysis

We assemble a novel dataset that measures the share of digitally-produced goods in the consumption baskets of households between 1960-2017. Our approach combines industry-level capital stocks with household-level consumption data and proceeds along the following steps: First, we

[1]See Veldkamp and Chung (2019) for a conceptual discussion about the measurement of data and digital technology content in the aggregate economy.

create an industry-level measure of ICT use by computing the share of ICT capital in an industry's value-added. Then, we trace linkages across industries via input-output tables to create ICT shares at the level of final goods and services. We match final goods and services to consumption product categories and calculate the digital share of consumption baskets along the income distribution. Our data reveal that rich households consume goods that are more ICT-intensive than poor households. Adding price data, we document that prices of ICT-intensive goods have grown at a slower pace than prices of non-ICT goods, making the rich households the prime beneficiaries of digitalization. We show how a significant part of observed consumption price inequality in the long-run can be explained by differences in the ICT content of the consumption basket. To cover consumption categories without market price, such as freely available goods or home production, Appendix A.6 links our digitalization measure to time-use data. We document that richer households spend more time on digital-intensive activities, confirming the pattern evident from the expenditure data. We also discuss alternative ways of measuring digitalization.

2.1 Capital Prices and Returns

The BEA's fixed asset accounts contain industry and aggregate-level details about the use of 96 different types of capital. We group these assets into two types of capital: ICT and non-ICT capital. We focus on equipment and intellectual property products. ICT capital includes the following asset categories: Mainframes, PCs, DASDs, printers, terminals, tape drives, storage devices, system integrators and intellectual property products, such as: prepackaged software, custom software, own-account software, semiconductor and other component manufacturing, and computers and peripheral equipment manufacturing. All other equipment and intellectual property assets as well as structures are defined as non-ICT capital. Appendix Tables A1 and A2 provide a detailed list of ICT and non-ICT assets. Our definition is closely connected to those of e.g. Eden and Gaggl (2018) and Karabarbounis and Neiman (2019), to which Section A.3 provides a comparison.

We retrieve capital prices from the BEA's National Income and Product Accounts (NIPA). Detailed information on this and the other datasets used in this paper can be found in Appendix Section A.2. Here and in the rest of the paper, we deflate nominal series and prices by a GDP deflator. This is because we are not interested in overall price inflation, but in price changes in one type of capital or product relative to others. Over the last decades, there have been substantial adjustments in the quality of capital, and new types of capital have been introduced. This applies in particular to digital capital. The BLS tries to take into account quality changes and the introduction of new products using methods such as hedonic quality adjustments and direct price comparisons.[2] However, this adjustment is not sufficient for ICT capital. Therefore, we follow the procedure suggested by Lashkari et al. (2024) and apply an additional adjustment term calculated by Byrne and Corrado (2017a) and Byrne and Corrado (2017b). These papers have calculated price adjustments by collecting data on technical improvements at the detailed product-level over time. For example, they include data on the amount of memory, clockspeed and other performance measures. They also provide more fine-grained software price estimates than the BEA by separately measuring prices of custom components and prepackaged software products.

[2]Detailed information on the methodology underlying the construction of these estimates is available at https://apps.bea.gov/national/pdf/Fixed_Assets_1925_97.pdf.

We aggregate the asset-level series into ICT and non-ICT prices using a Törnqvist index. The top part of Figure 1 shows the resulting price series for ICT and non-ICT capital, $p_{ICT,t}$ and $p_{NICT,t}$, between 1960 and 2017. While the price of non-ICT capital has seen a modest increase, the price of ICT capital has decreased massively at a rate of about 10% per year. By 2017, the price was only about 1 % of its 1960 level. Thus, it has become much cheaper for industries to use ICT capital over time.

We calculate the rate of return on ICT and non-ICT capital following Karabarbounis and Neiman (2019) as

$$R_{k,t} = p_{k,t}\frac{(1+\tau_t^x)}{1-\tau_t^k}\left[\left(\frac{(1+\tau_{t-1}^x)p_{k,t-1}}{(1+\tau_t^x)p_{k,t}}\right)(1+(1-\tau_t^k)\tilde{r}_t) - (1-\delta_{k,t}) - \frac{\tau_t^k\delta_{k,t}}{1+\tau_t^x}\right] \tag{1}$$

for $k = ICT, NICT$. $p_{k,t}$ is the (deflated) price of capital, $\delta_{k,t}$ the depreciation rate, τ_t^x the tax rate on investment, τ_t^k the tax rate on capital and $\tilde{r}_t$ the cost of capital. The cost of capital consists of two elements. First, the real interest rate, measured as the nominal rate on 10-year US Treasuries minus a five-year moving average of realized inflation that proxies expected inflation. Second, and differently from Karabarbounis and Neiman (2019), we add a constant 5% risk premium to to reflect that the costs of borrowing for firms are higher than the risk-free rate. This follows the approach of Barkai (2020) for equity capital and results in costs of capital that closely correspond to his measure (see Appendix Figure A1). Note that eq.(1) is an approximation derived from the optimality conditions of a representative firm and may not necessarily reflect the true user cost of capital.

The bottom part of Figure 1 shows the rate of return on ICT capital over the sample period. We plot the real return, i.e. $\frac{R_{k,t}}{p_{k,t}}$ as the nominal rate for ICT is dominated by the extreme decline in the capital price. The real return is also the more meaningful economic measure. As the figure shows, the real return on ICT capital is always above the real return on non-ICT capital at between 40-60% per year. Next to the larger price decline, this is due to a higher depreciation of ICT capital.

2.2 Digitalization at the Industry Level

We combine the capital price and returns series with industry-level data on asset use as well as capital stocks and value-added, using the BEA's Detailed Data on Fixed Assets and Industry GDP data. We calculate the **ICT share** for industry j at time t as

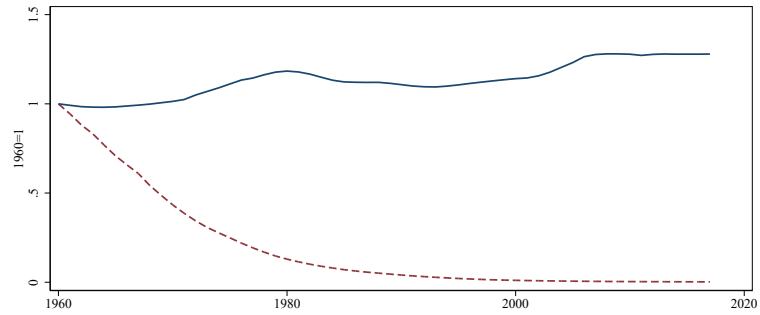
$$sK_{ICT,t}^j = \frac{R_{ICT,t}K_{ICT,t}^j}{P_t^j Y_t^j} \tag{2}$$

where $K_{ICT,t}^j$ is the real capital stock and $P_t^j Y_t^j$ is value-added. This definition of digitalization reflects the contribution of ICT capital to value-added and therefore captures the relative importance of ICT in the production process.

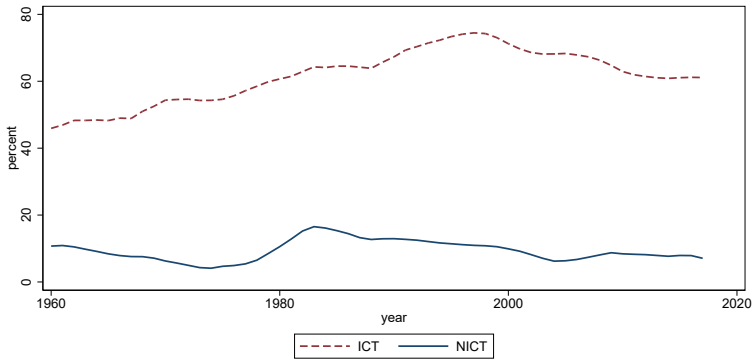
Figure 2 (a) depicts the ICT share over time. We show the mean, the median and the industries between the 10th and 90th percentile. The average ICT capital share has risen from almost zero

Figure 1: ICT- and non-ICT capital prices and returns

(a) prices $p_{k,t}$



(b) real returns $R_{k,t}/p_{k,t}$



Note: The top figure shows five-year moving averages of the prices of ICT and non-ICT capital, normalized to 1 in 1960 and adjusted for CPI inflation. The bottom figure shows real rates of return on ICT and non-ICT capital. *Source:* BEA and own calculations.

in 1960 to 5.1% in 2017. This number is comparable to numbers reported by Karabarbounis and Neiman (2019) and Eden and Gaggl (2018) (see Appendix A.3). The ICT share exhibits a large degree of heterogeneity across industries. Industries in the top 10% of the ICT share mostly belong to the computer and electronics manufacturing industry or the information industry. However, this group also includes some financial services and management. Among the least ICT-intensive industries are agriculture, construction and mining.

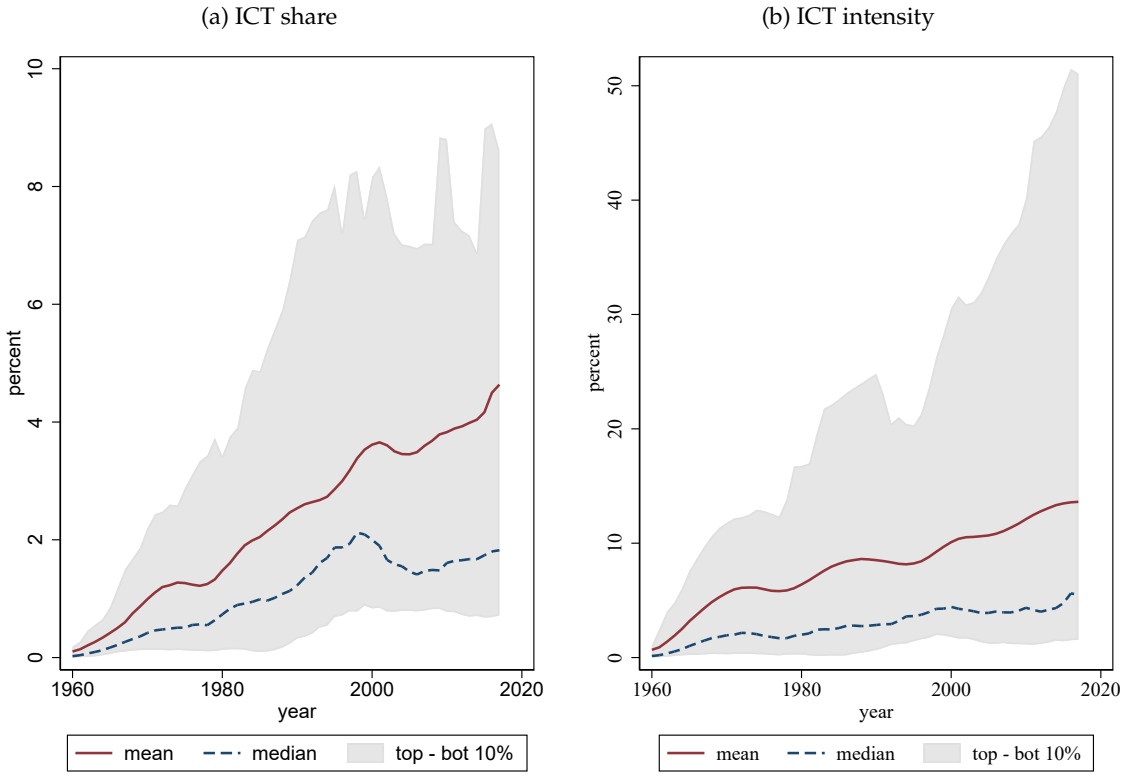
As discussed above, the rate of return used for the ICT share estimation is not directly observed. An alternative measure, which only considers well-measured objects – capital prices and quantities – is the share of ICT capital in the total capital stock. Figure 2 (b) shows this ICT intensity over time.[3] The share of ICT capital in the total capital stock increased on average from less than 1% in 1960 to over 13% in 2017. For some industries, it even increased to more than 50% of capital inputs.

Note that both measures include two counteracting trends: On the one hand, a massive decline in the price of ICT capital, as shown in Figure 1 (a). On the other hand, a strong increase in the relative use of ICT capital. Nevertheless, both indicate an increased importance of digitalization.

Industries may not only use digital capital in their own production activities, but also use intermediate inputs that have been produced with digital capital. In order to calculate the share of digital capital used in the production of final goods and services, we need to take input-output linkages

[3] In a previous version of the paper, we used this as our main measure of digitalization.

Figure 2: Industry-level digitalization

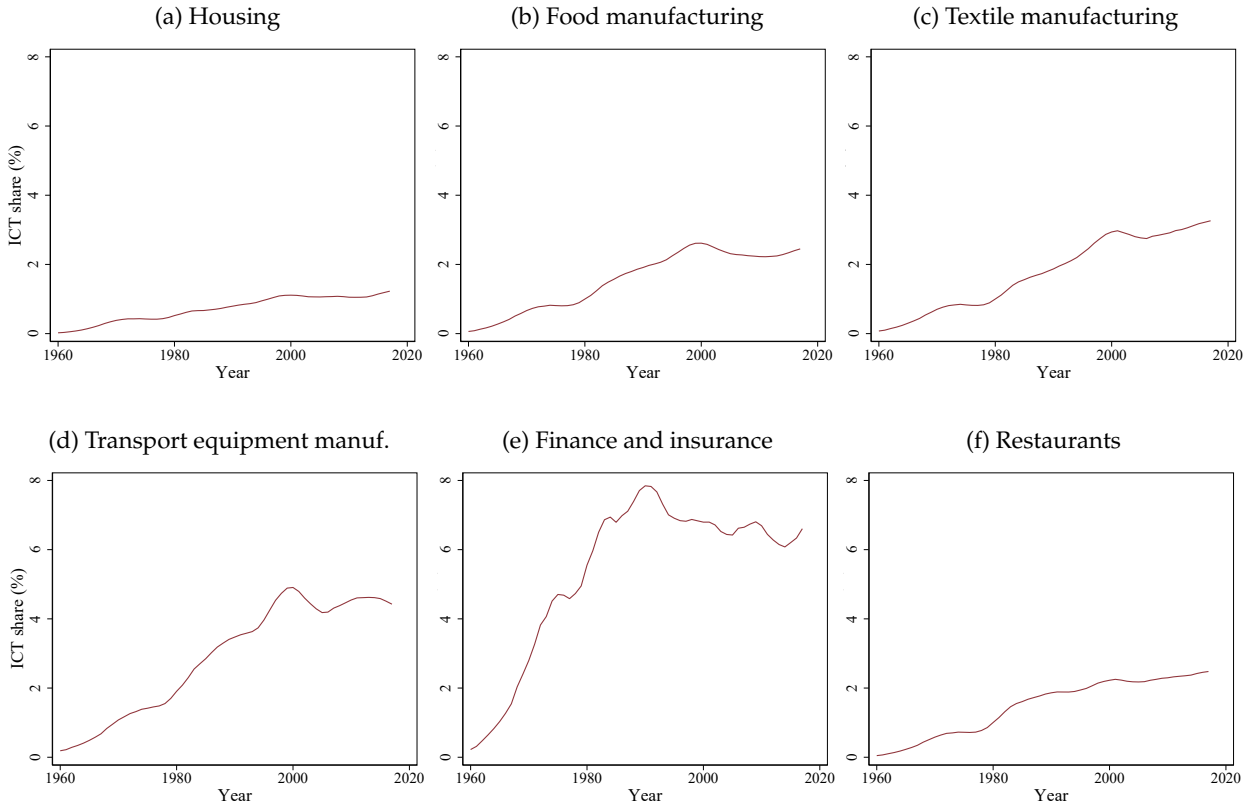


Note: The left graph plots the share of ICT in value-added. The right graph plots the ICT intensity, defined as the share of ICT capital in the total capital stock. The solid line shows the mean, the dashed line the median and the gray area the industries between the 10th and 90th percentile of the distribution. *Source:* BEA and own calculations.

between industries into account. The BEA's Input Output (IO) Accounts show how industries provide input to or use output from each other. We use the detailed IO tables after re-definitions. The tables are published every five years and contain around 400 private industries, which can be matched to the 61 BEA industries. We create a commodity-by-commodity direct requirements matrix, which we use iteratively to calculate the digital share of final goods and services as a weighted mean of the digital shares of its intermediate inputs and value added. Goods and services that initially have a low degree of digitalization tend to be more digitalized after the inputs from other industries are considered. The reverse is true for highly digitalized industries. Appendix A.2 provides further details on the IO tables and our procedure.

We compute the ICT share for all final goods and services between 1960 and 2017. Grouping these into 25 broader product categories, Figure 3 plots the ICT share of six aggregate categories over time. In anticipation of the link to consumption categories in the CEX in Section 2.3, we consider the six most important categories for consumption: Housing, food manufacturing, textile manufacturing, transport, finance and insurance and restaurants. There has been an increase in the ICT share in all of the product categories over time and the pattern often looks similar, e.g. reflecting the build-up and subsequent burst of the dotcom bubble. However, product categories are characterized by large differences in the average ICT share. In particular, finance and insurance has a high ICT share, whereas food manufacturing is characterized by a low ICT share. These differences will become relevant in the construction of ICT shares for individual consumption baskets.

Figure 3: ICT share by product category



Note: The graph shows ICT shares for six broad product categories over time. The ICT share in these categories is calculated as a weighted average of ICT share of the relevant goods or services. The weights correspond to their share in value added. *Source:* BEA and own calculations.

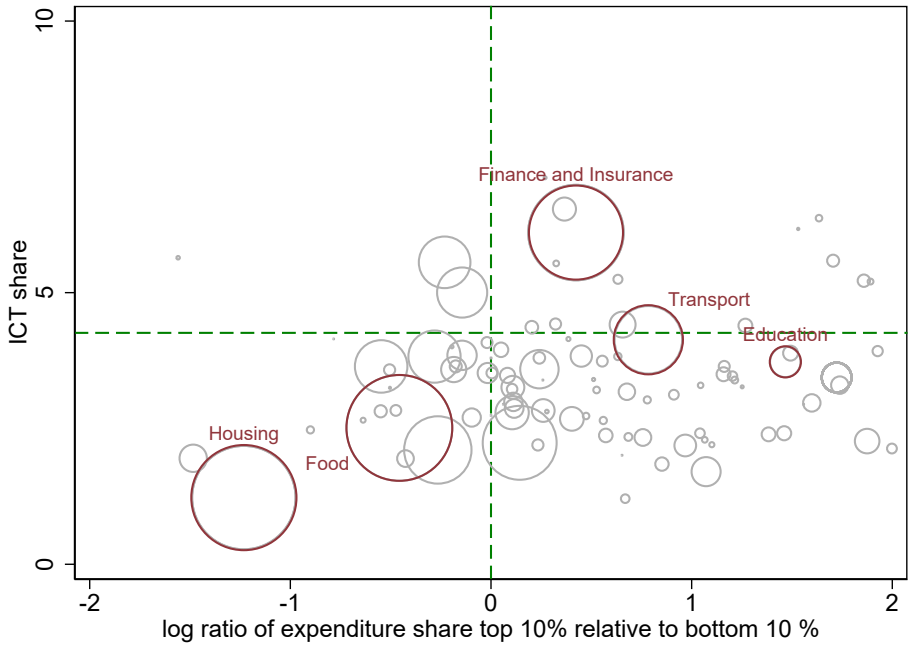
2.3 Digitalization and Consumption

We calculate the digital share of household consumption along the income distribution by matching final goods and services to expenditure categories from the Consumer Expenditure Survey (CEX). The CEX is the most detailed expenditure survey in the United States, carried out at the household level. Next to data about the purchases of hundreds of disaggregate goods and services, it also contains a large amount of demographic and income information on the households. We use the interview survey from the public-use microdata for 1984-2017. Between 1990-2017, data are available at the detailed product level. Between 1984-89, we follow Jaravel and Lashkari (2024) in imputing detailed expenditures based on more aggregate expenditure categories. More details on the survey and on the construction of the dataset can be found in Appendix A.2. We divide households into equal-sized bins based on total household income. For each income group, we create a weighted average of expenditure by year. Based on the concordance table of Borusyak and Jaravel (2021), we link the CEX product categories to IO product categories. Each product code in the CEX is matched to a unique IO code.[4]

Income groups vary in their spending pattern, and the relative consumption of different products

[4]Note that the concordance table treats housing as a service from the real estate sector rather than as a good from the construction sector, in line with our procedure when matching BEA industries to IO tables. Housing is also special since it serves both as consumption and investment good for house owners. We take this into account by using the rental value of owner-occupied housing, which equals its consumption value.

Figure 4: ICT share vs. relative expenditure shares by product category in 2017



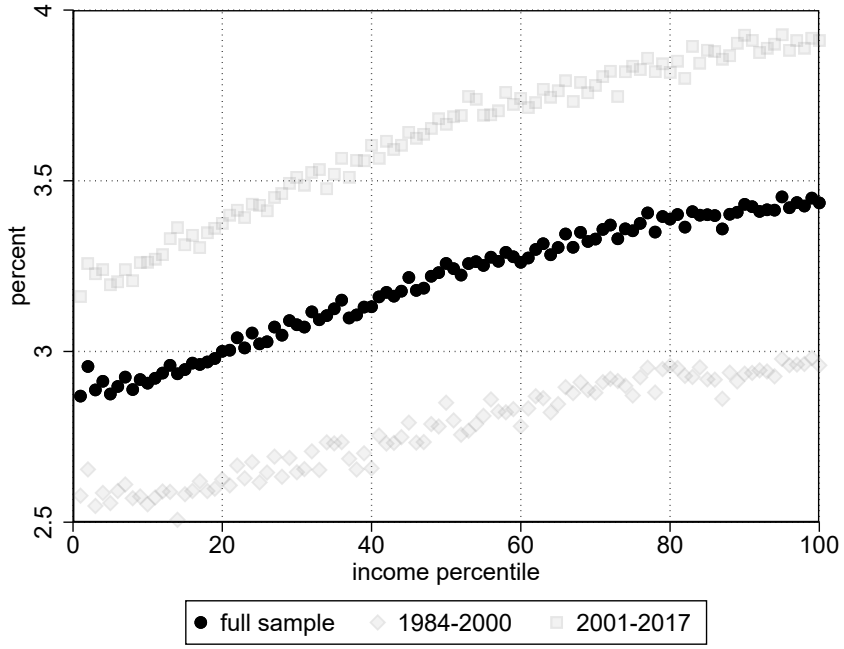
Note: Each circle corresponds to an IO product category. The size of the circle reflects the share of expenditures for that category of the median household. The x-axis shows the log ratio of the expenditure share of the top 10 % relative to the bottom 10 %. The more to the right the data are, the more important that category for the rich. Values around zero indicate that a category is equally important to the top 10% as it as for the bottom 10 %. The y-axis shows ICT share of the industry, the dashed green lines indicates average ICT share and equal expenditure shares, respectively. Data are for 2017, the graph is zoomed into expenditure categories with an ICT share below 10% and a cost share ratio below 2. *Source:* BEA, CEX and own calculations.

by rich and poor households are systematically correlated with the ICT share of these goods. Figure 4 plots the ICT share of goods and services against its relative importance for high-income vs. low-income households in 2017. The pattern is very similar for earlier years. Each dot corresponds to an IO product category and the size reflects the importance of the product (in terms of expenditure shares) for the median household. For clarity of exposition, we only show categories with an ICT share below 10 and with a cost share ratio below 2.[5] The dashed green line represents the average ICT share. In the negative area of the x-axis, where categories are more important for poorer households, most ICT shares are below the average value. In this category are in particular basic goods like housing and food. In the positive area of the x-axis which depicts consumption goods important for the rich are some categories with very large ICT shares such as finance and insurance. Overall, there is a positive correlation between the ICT share and the relative expenditure share of high-income households.

In the final step, we compute the ICT share of the whole consumption basket by income percentile as an expenditure-weighted average ICT share of the individual product categories. Figure 5 shows the resulting measure. We show both averages across the whole sample period and separately for 1984-2000 and 2001-2017. The ICT share increases monotonically along the income

[5] 11 categories relating to computer hardware and software have an ICT share between 21 and 36. These are mostly consumed by richer households.

Figure 5: ICT share along the income distribution



Note: The graphs show the ICT share of consumption along the income distribution. The black markers show averages across the whole time period 1984-2017, the gray markers show averages for subperiods. *Source:* BEA, CEX and own calculations.

distribution. Throughout the sample period, there is a 21% (0.57 percentage points) difference between the top and bottom. The slope is slightly steeper for the bottom half of the distribution, flattening out for higher income percentiles. Comparing the two sub-periods reveals a strong increase in the ICT share over time. The median increased from 2.85 in the first half of the sample to 3.67 in the second half. At the same time, the absolute difference between top and bottom has increased, with the result that the income gradient has stayed approximately constant.

In Appendix A.4, we carry out several robustness checks of our results. First, we create percentiles of labor income instead of total income, which excludes people at the bottom of the distribution who receive government support and do not work. The resulting graph still looks very similar. Second, we exclude finance and insurance, which may be considered an investment good rather than a consumption good. Excluding this category lowers the ICT share by about 0.2 percentage points for all households. Third, we exclude durable goods. We still see that the ICT share increases along the income distribution, but the gradient becomes smaller. Fourth, we show figures based on historical consumption data for the 1960s and 1970s, based on the dataset by Jaravel and Lashkari (2024). The authors use summary expenditure data for 1960, 1972 and 1980 to impute detailed expenditure at the annual level. Since this procedure introduces considerable noise, we do not include the 1960-84 data in the main analysis. However, we still observe a positive income gradient during the 1960's and 1970's. Finally, we show that the results continue to hold when controlling for demographic factors. In panel regressions, we regress the ICT share on income, age, education and region and show that even though higher-income households tend to be older and more educated, income is still a significant predictor of the ICT share. This is also shown in Figure A9, which replicates Figure 5 for demographic subgroups.

2.4 Digitalization and Consumer Prices

To what extent does the heterogeneity in ICT shares and the large decline in ICT input prices translate into lower relative consumer prices? In principle, the price of a good or service is determined by its costs of production and the profit margin. Production costs include labor, capital and intermediate inputs. Abstracting from profits, changes in prices can be predicted by changes in prices of each input factor, multiplied by its value-added share. Muting all other factors, the effect of digitalization can be approximated by

$$\Delta \hat{p}^{j}_{t,t-s} = sK^{j}_{ICT,t} \Delta p_{ICT,t,t-s} \tag{3}$$

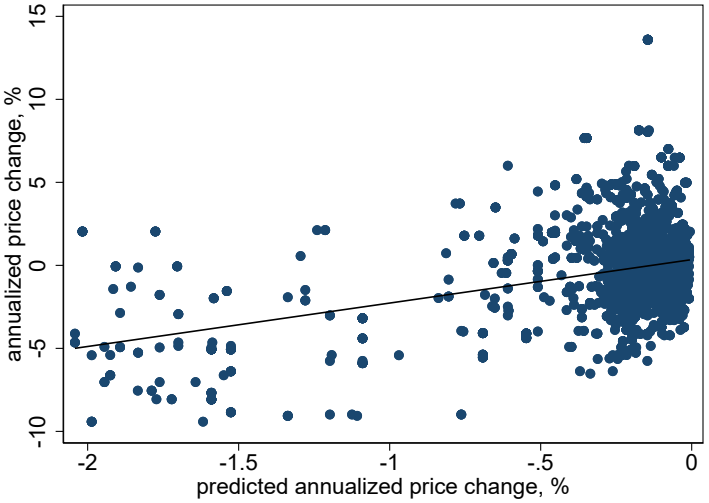
where, as before, $sK^{j}_{ICT,t}$ is the share of ICT capital in value-added of industry j and $p_{ICT,t}$ the price of ICT capital. This equation delivers a prediction of the change in output prices in a scenario where only ICT input prices change. It abstracts from increasing returns to scale properties of digitalization in the production process that is discussed by Lashkari et al. (2024). Comparing the prediction of equation (3) with actual price changes nevertheless provides a rough lower bound assessment of what share of actual price changes digitalization can explain.

We use NIPA data on output prices and match each of the different expenditure categories to an IO product category. Price data by NIPA aim to take quality changes in products over time into account (NIPA, 2023), which is particularly relevant for digitally-produced products. In particular, hedonic adjustments aim to decompose items into their characteristics and attach values to each of these characteristics. Consumption price changes related to computers or other ICT goods thus take into account that, for example, PCs have gotten more memory. However, this adjustment is not undertaken in real time, it does not extend to all consumption categories and neither systematically accounts for product entry and exit. While new product bias should disappear over long horizons, unmeasured quality improvements will lead to a bias in the price data. If these adjustments are larger for ICT products than for non-ICT products, the results shown in this section probably understate the full effect of ICT on prices.

We deflate all price changes with CPI inflation and take long rates of changes in prices across 10 year intervals to smooth out business cycle fluctuations. Figure 6 plots actual price changes for each final product and decade against price changes predicted by eq 3. There is a significant positive correlation. However, in the figure we do not control for other relevant drivers of price changes. Some industries such as oil-producing and petrol refining industries, have seen strong price increases due to external shocks. Others, such as the tobacco industry, health or education, have seen price increases due to changes in taxation and subsidies. Additionally, if an industry's ICT share is correlated with its import share, Figure 6 might erroneously pick up trade integration, which led to price declines in many consumer products over the last decades.

Table 1 therefore shows outcomes of stacked regressions of annualized decadal relative price changes on the ICT share. The regressions are at the level of product categories. We always include fixed effects at the level of 16 major industries, corresponding to NAICS sectors. In columns (2) and (4), we also include decade fixed effects. Additionally, in columns (3) and (4) we control for the share of imports in an industry's value-added. Since import data at the detailed industry level

Figure 6: Final product price changes and ICT 1960-2017



Note: The figure plots average annual price changes during a decade against predicted price changes at the level of expenditure categories. *Source:* BEA and own calculations.

are not available for the whole period, these regressions cover only 1997-2017. In all regressions, we cluster standard errors at the level of IO industries.

The coefficient on the ICT share is always negative and significant. In column (1), a 1 percentage point increase in the ICT share coincides with a 0.17 percentage points (17 basis points) lower annual increase in the relative price. The coefficient moves closer to zero when including additional controls, but remains negative and highly significant throughout all specifications. In Appendix Table A.5, we carry out further extensions and robustness checks. We show that the relationship between ICT share and price changes is nonlinear, as a quadratic term is highly significant. Furthermore, we show that the ICT share is also significant when measured at the beginning of the ten-year period instead of at the end.

While these results should not be interpreted causally, they nevertheless provide compelling evidence that prices of consumption goods co-move negatively with the good's ICT share, even after controlling for potential confounders. Our findings align with Aghion et al. (2020), who show that automation reduces the producer price index in French manufacturing industries. Graetz and Michaels (2018) find that the use of robots – another type of automation technology – also leads to lower prices. Comparing the magnitudes of the two approaches, the regression approach results in slightly stronger ICT-driven price changes than the direct ICT price approach from eq. (3): Going from the 10th to the 90th percentile of the ICT share distribution would lead to a decline in price levels by -0.483 percentage points according to the regression approach (column (6) of Table 1) and -0.323 percentage points according to the direct approach.

Table 1: Price changes and ICT share

Variables	(1) 60-17	(2) 60-17	(3) 60-17	(4) 97-17	(5) 97-17	(6) 97-17
ICT share	-0.166***	-0.137***	-0.148***	-0.147***	-0.107***	-0.107***
	(0.0488)	(0.0308)	(0.0356)	(0.0374)	(0.0293)	(0.0293)
trade share				-0.000348***	-0.000267***	-0.000268***
				(0.000118)	(8.93e-05)	(9.13e-05)
Constant	0.328	0.226	0.111	0.285	0.331	0.0606
	(0.274)	(0.203)	(0.0875)	(0.320)	(0.207)	(0.185)
Observations	3,912	3,912	3,912	1,160	1,160	1,160
R-squared	0.102	0.343	0.346	0.209	0.358	0.358
industry FE	no	yes	yes	no	yes	yes
decade FE	no	no	yes	no	no	yes

Note: The table shows the regression of annualized decadal price changes on the ICT share and the import share at the end of each decade. The unit of observation is the product category. There are a total of 674 product categories, with some missing values for early years. All variables are measured in %. Columns (1), (2) and (3) are regressions for the full dataset, columns (4), (5) and (6) from 1997. Standard errors clustered at the level of final goods and services in parentheses. *** p<0.01, ** p<0.05, * p<0.1

2.5 Diverging Costs of Living

Combining our findings on price differences of product categories from Section 2.4 with the different consumption baskets derived in Section 2.3, we now consider changes in consumer prices along the income distribution. While it has been documented that poorer households have experienced systematically higher inflation over the last decades (Jaravel, 2019; Kaplan and Schulhofer-Wohl, 2017), we study what share of this difference can be explained by differences in the ICT share.

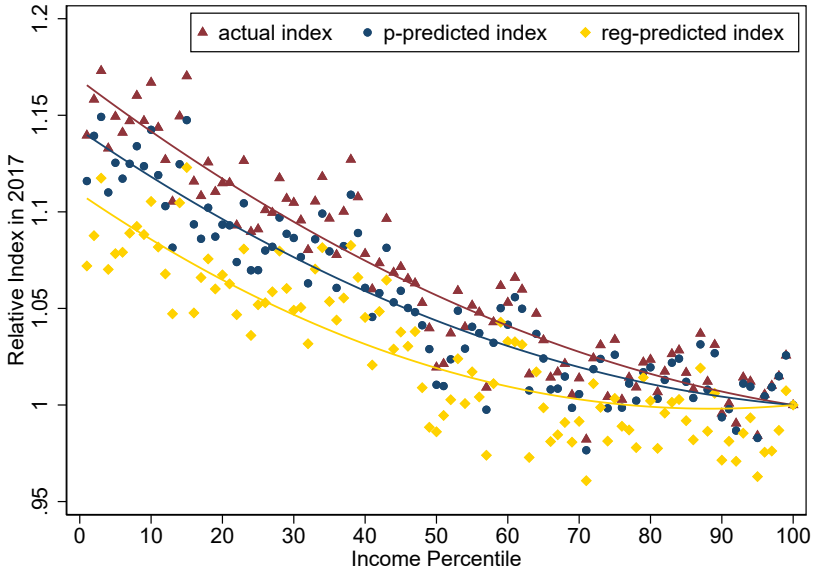
To construct relative price changes across income percentiles, we follow Jaravel (2021) in calculating Törnqvist price indices at the percentile level. This implicitly assumes that preferences within each income percentile are homothetic. Price changes are expressed as

$$1 + \pi^i_{t,t+1} = exp\left(\sum_{j=1}^{n} \frac{s^i_{j,t} + s^i_{j,t+1}}{2} \log\left(\frac{p^i_{j,t+1}}{p^i_{j,t}}\right)\right)$$

where i refers to the income percentile, j to the product and t to the time period. $s^i_{j,t}$ and $s^i_{j,t+1}$ are expenditure shares in periods t and $t+1$, respectively. The Törnqvist index is a natural choice as it is a superlative index that weights products in the consumption basket by an average of period t and t+1 shares, thus taking into account product substitution. We show the robustness of our results to the use of a Laspeyres and Paasche price indices in Appendix Figure A10.

We use NIPA data as in Section 2.4 and merge these with our consumption dataset from the CEX to calculate changes in consumer prices for each income percentile between 1984 and 2017. The red triangles in Figure 7 show the increase in the percentile-specific price index between 1984 and 2017. The price index is deflated by the price change in the highest income percentile, such that all data points are relative to percentile 100. Consumer prices increased most strongly at the bottom

Figure 7: Relative consumption prices and ICT



Note: The graph plots changes in the price of consumption across the income distribution between 1984 and 2017, where the price changes in the highest income percentile are normalized to 1. The red triangles indicate the actual data, whereas the blue circles and the yellow diamonds indicate price indices in the absence of digitalization. Blue: prices predicted following the approach of Figure 6. Yellow: prices predicted by regression coefficients of Table 1. The markers represent the data points and the curves estimated polynomials of degree 2. *Source:* BEA, CEX and own calculations.

of the income distribution, where the price index is about 15 percent higher than at the top. It flattens out for the top 30% of the distribution. Overall, the differences suggest a substantial shift in relative purchasing power along the income distribution.

To isolate the role of digitalization for these price changes, we plot the actual price changes against two alternative price indices, which hypothesize that digitalization had not taken place. The first price index follows the logic of Section 2.4 and predicts ICT-driven price changes by using the ICT share of a product together with the evolution of the price of ICT capital. We use $\Delta\hat{p}^{j}_{t,t-s}$ together with the percentile-specific expenditure shares of each good j, $s^{i}_{j,t}$. The blue circles in Figure 7 show the difference between actual and ICT-predicted price index, i.e. they show prices in the absence of digitalization. We have again normalized price changes in percentile 100 to 1. According to the blue fitted polynomial shows, price changes at the bottom would have been about 3 percentage points lower if digitalization had not taken place. The digitalization effect therefore accounts for about one fifth of the total relative price difference.

We compute a second hypothetical price index based on the regressions shown in Table 1 and predict ICT-driven price changes by the estimated coefficient multiplied with the average ICT share of each income percentile. We plot again the difference between actual and predicted price index. As the golden diamonds in Figure 7 show, this approach implies a larger role of digitalization for explaining differences in relative prices across the income distribution. For the first percentile, the price index under the no-digitalization scenario is about 7 percentage points lower in 2017 than the actual price index. This suggests digitalization accounts for about two fifth of the total relative price difference.

How much do relative price changes matter for welfare? This question cannot be answered based

on data alone but requires a specification of preferences. As shown by Baqaee and Burstein (2023), assuming homothetic preferences within income percentiles may lead to biases in welfare calculations, as it does not properly take expenditure switching due to income effects into account. Therefore, in Section 3, we introduce a model that specifies the preference structure and allows us to make welfare statements.

2.6 Discussion

We have calculated the ICT share of final goods by whether they are produced with digital capital. This follows the literature on the employment effects of digitalization, which treats digital technology as an input in the production process. By using the same approach, we can directly compare the income and the price effect of digitalization. However, it is important to note that this is not the only possible way of calculating ICT shares. An alternative measure would classify goods based on whether they themselves constitute digital goods (BEA, 2021). By this output-based classification, a smartphone would always be a digital good, whereas by our input-based definition, it would constitute a digital good only if it is produced with digital capital.

Appendix A.6 reproduces Figure 5 using the output-based approach. The ICT share now exhibits a hump-shape instead of a monotonic increase. It is therefore important to bear in mind that any results documented in the previous sections are contingent on the definitions applied. A similar argument is made in the structural change literature, which considers two different approaches for defining sectors, one based on value-added, and one based on final consumption purchases by the households (Herrendorf et al., 2013). Depending on the question of interest, either one might be more suitable.

A second discussion point concerns the use of expenditure data for measuring the ICT share of consumption. This approach has two shortcomings. First, some households might substitute market expenditures with home production (Aguiar and Hurst, 2005). Second, there are digital goods and services that are available for free, or at a very low marginal cost (think about streaming services or online dictionaries). These are not properly accounted for in the expenditure data. To address these issues, in Appendix A.6 we supplement our analysis with data from the American Time Use Survey (ATUS), which documents how Americans allocate their time across different activities. Analogous to the ICT share of expenditures in Figure 5, Figure A13 (a) shows the ICT share of time use along the income distribution. This ICT share, while more noisy than the expenditure-based measure, also exhibits an upward-slope, and is therefore in line with the expenditure-based measure. Furthermore, zero marginal cost products are not consumed more intensively by poor households, as indicated by Figure A13 (b). There is also no systematic evidence for substitution of market expenditures with home production for ICT intensive goods, as there is a near zero correlation between expenditure and time-shares across income groups. In conclusion, focusing on expenditure data only should not result in a bias.

3 Model

In the following, we model the effect of an increase in digital technology on consumption inequality, working through changes in relative prices and income. In the model economy, there are two capital goods (ICT and non-ICT) and two sectors with different factor inputs. Each sector produces a unique final consumption good. Workers have non-homothetic preferences over the two goods, which means that income differences induce differences in consumption patterns. There are two types of workers, high-skill and low-skill, where low-skill workers sort into routine and manual occupations based on individual ability. Technological advances affect the demand for factor inputs and worker incomes. This setting allows us to quantify the effect of digitalization on consumption inequality comprehensively and to analyze the mechanisms at work.

3.1 Sectoral production functions

Two sectors $i = 1, 2$ each produce a final output using as inputs two types of capital, ICT capital $K_{ICT,i}$ and non-ICT capital, $K_{NICT,i}$ and three types of labor: (Low-skill) routine labor R_i, (low-skill) manual labor M_i and (high-skill) cognitive labor H_i. Production is nested.

Inner nest: Routine labor R_i and ICT capital $K_{ICT,i}$ are used to produce an automated work output (AW):

$$AW_i = \left[\phi_i R_i^{\frac{\eta_i - 1}{\eta_i}} + (1 - \phi_i) K_{ICT,i}^{\frac{\eta_i - 1}{\eta_i}} \right]^{\frac{\eta_i}{\eta_i - 1}}, \tag{4}$$

with $\phi_i \in (0, 1)$ weights that govern the importance of each input factor, and η_i the elasticity of substitution between routine labor and ICT. If η_i is large, then routine labor tends to be more substitutable to ICT. If $\eta_i = 1$, the production function is Cobb-Douglas with factor share ϕ_i.

Middle nest 1: High-skill labor H_i and the automated work output are used to produce an intermediate product SW:

$$SW_i = \left[\gamma_i H_i^{\frac{\epsilon_i - 1}{\epsilon_i}} + (1 - \gamma_i) AW_i^{\frac{\epsilon_i - 1}{\epsilon_i}} \right]^{\frac{\epsilon_i}{\epsilon_i - 1}}, \tag{5}$$

with weights $\gamma_i \in (0, 1)$ and ϵ_i the elasticity of substitution between high-skill labor and AW. Smaller values indicate that both inputs are complements, larger values that they are substitutes.

Middle nest 2: SW and manual labor are combined to produce total work output (TW) in a Cobb-Douglas manner with factor share ψ_i:

$$TW_i = M_i^{\psi_i} SW_i^{1 - \psi_i}$$

which is motivated by the observation of Eden and Gaggl (2018) and Jaimovich et al. (2021) that the share of non-routine manual labor in national income has stayed fairly constant over time.

Outer nest: Total work and non-ICT capital are combined in a Cobb-Douglas production function,

owing to the fact that the share of non-ICT capital in revenue has been constant over the last decades (Koh et al., 2020). Each sector also has a specific level of Total Factor Productivity (TFP) A_i at time t,

$$Y_i = A_i K_{NICT,i}^{\alpha_i} T W_i^{1-\alpha_i} \tag{6}$$

There is perfect competition in each sector. Firms demand a price P_i for their good. They pay wages W_H, W_R and W_M and rent capital which pays interest rates R_{NICT} and R_{ICT}. First-order conditions yield the following equations for factor returns

$$K_{NICT,i}: \quad P_i \alpha_i \frac{Y_i}{K_{NICT,i}} = R_{NICT} \tag{7}$$

$$M_i \quad P_i (1-\alpha_i) \psi_i \frac{Y_i}{M_i} = W_M \tag{8}$$

$$H_i: \quad P_i \gamma_i (1-\alpha_i)(1-\psi_i) \frac{Y_i}{SW_i^{\frac{\epsilon_i - 1}{\epsilon_i}} H_i^{\frac{1}{\epsilon_i}}} = W_H \tag{9}$$

$$R_i: \quad \frac{P_i \cdot AW_i^{\frac{\epsilon_i - 1}{\epsilon_i}} (1-\gamma_i)(1-\psi_i)(1-\alpha_i)\phi_i Y_i}{R_i^{\frac{1}{\eta_i}} AW_i^{\frac{\eta_i - 1}{\eta_i}} SW_i^{\frac{\epsilon_i - 1}{\epsilon_i}}} = W_R \tag{10}$$

$$K_{ICT,i}: \frac{P_i \cdot AW_i^{\frac{\epsilon_i - 1}{\epsilon_i}} (1-\gamma_i)(1-\psi_i)(1-\alpha_i)(1-\phi_i) Y_i}{K_{ICT,i}^{\frac{1}{\eta_i}} AW_i^{\frac{\eta_i - 1}{\eta_i}} SW_i^{\frac{\epsilon_i - 1}{\epsilon_i}}} = R_{ICT} \tag{11}$$

3.2 Technological progress and capital formation

The output of sector 1 can be used for consumption as well as to produce non-ICT capital while the output of sector 2 can be used for consumption and for the production of ICT capital. We define μ_{ICT} as the rate of transformation of Y_2 into ICT capital and μ_{NICT} as the rate of transformation of Y_1 into non-ICT capital.[6] Then, the resource constraints are

$$Y_1 = C_1 + \mu_{NICT} I_{NICT}, \qquad Y_2 = C_2 + \mu_{ICT} I_{ICT}. \tag{12}$$

μ_{ICT} is the relative price of ICT capital and at the same time measures exogenous progress in ICT technology (see e.g. Greenwood et al., 1997; Karabarbounis and Neiman, 2014; Eden and Gaggl, 2018; Cortes et al., 2017 for similar set-ups). Similarly, μ_{ICT} measures improvements in non-ICT

[6] We discuss this modeling choice and the calibration further in the Section 4.1 when we assign sectors to their empirical counterparts.

technology. A decline in μ_{ICT} or μ_{NICT} makes production of the respective type of capital cheaper[7]. This will translate into lower production costs, in particular in the ICT-intensive sector 2, which will be passed on to consumers in form of lower prices.

In addition to this, we consider sector-specific TFP growth, which contributes to the increase in production as discussed by Greenwood et al. (1997). The relative magnitude of this source of growth will be calibrated to match key moments in the data for the main benchmark model. TFP A_i grows at an exogenous and constant annual rate $g_{A,i}$

$$A_i' = (1 + g_{A,i})A_i \tag{13}$$

Capital follows the standard law of motion

$$K_{NICT}' = (1 - \delta_{NICT})K_{NICT} + I_{NICT}, \qquad K_{ICT}' = (1 - \delta_{ICT})K_{ICT} + I_{ICT}, \tag{14}$$

where δ_{NICT} and δ_{ICT} are the depreciation rates.

3.3 Workers

Occupational Choice

There are two types of workers in this economy, high-skill and low-skill. The worker groups are of sizes $\bar{H}$ and $\bar{L}$, which are time-varying. We denote by g_H and g_L the exogenous growth rates of the two groups. We think about skill in terms of level of education, which is determined before entering the labor market. Workers can work in any sector and can switch sectors at no cost. Following Jaimovich et al. (2021), we assume that high-skill workers are all identical and always allocate to the high-skill tasks, whereas low-skill workers may work in either routine or manual occupations. Each low-skill worker j is endowed with two idiosyncratic productivity parameters $\lambda_{R,j}$,$\lambda_{M,j}$ indicating their routine and manual abilities, which are drawn from a joint distribution $\Gamma(\lambda_R, \lambda_M)$. More able workers supply more units of labor and therefore earn a higher wage. A worker chooses to work in R if

$$\lambda_{R,j} W_R \geq \lambda_{M,j} W_M \tag{15}$$

Thus, occupational choice depends on inherent worker characteristics as well as on equilibrium wages. We assume that workers are free to switch occupation in response to changes in relative wages.

In consequence, total supply of routine labor (in efficiency units) is

$$\bar{R} = \bar{L} \int_{-\infty}^{\infty} \int_{-\infty}^{\lambda_M(\lambda_R)} \lambda_R \Gamma'(\lambda_R, \lambda_M) d\lambda_M d\lambda_R \tag{16}$$

where $\lambda_M(\lambda_R)$ denotes the cutoff in ability for each λ_R value such that, below it, workers choose

[7]Progress in ICT technology could be endogenized e.g. by adding an element of expanding varieties in digital capital or learning-by-doing as in Oberfield (2023).

to work in R and not in M. $\Gamma'(\lambda_R, \lambda_M)$ is the density function corresponding to the distribution function $\Gamma(\lambda_R, \lambda_M)$.

Correspondingly, the total number of manual efficiency units provided is

$$\bar{M} = \bar{L} \int_{-\infty}^{\infty} \int_{-\infty}^{\lambda_R(\lambda_M)} \lambda_M \Gamma'(\lambda_R, \lambda_M) d\lambda_R d\lambda_M \tag{17}$$

with $\lambda_R(\lambda_M)$ defined accordingly.

Preferences

All workers have non-homothetic preferences over the two goods 1 and 2. We follow Boppart (2014) by introducing a type of "price independent generalized linearity" (PIGL) preferences. Preferences are specified over the prices P_1 and P_2 and the nominal expenditure level of worker j, e_j. The indirect instantaneous utility function is

$$V(P_1, P_2, e_j) = \frac{1}{\rho} \left(\frac{e_j}{P_2} \right)^{\rho} - \frac{\nu}{\theta} \left(\frac{P_1}{P_2} \right)^{\theta} - \frac{1}{\rho} + \frac{\nu}{\theta}, \tag{18}$$

where the parameters govern how the demand for the two goods reacts to changes in relative prices (θ) and the income level (ρ). ν controls the average expenditure shares of goods 1 and 2. We do not make any explicit parameter restrictions at this point.[8] PIGL preferences offer several advantages. First, they generate non-homothetic behavior of individuals even as incomes become very large. This is relevant given the long time span we consider in our simulation and stands in contrast to, for example, generalized Stone-Geary preferences, which imply that the effect of the income level on consumption patterns tends towards zero for large income values. Second, PIGL preferences allow for easy aggregation. This is particularly relevant in the case of routine and manual workers with their heterogeneous abilities, as we will show below.

Optimization

As in Jaimovich et al. (2021), all non-ICT and ICT capital is owned by high-skill workers. This assumption is motivated by empirical evidence, as we show based on the Survey of Consumer Finances (SCF) in Table A3. Furthermore, we will show in our simulation that changes in capital income are not the driving force of increasing income inequality.

Since high-skill workers are all identical, we assume they form one household of size $\bar{H}$, which solves a joint optimization problem. The budget constraint for the high-skill household is

$$\underbrace{C_{1,H} P_1 + C_{2,H} P_2}_{e_H} + I_{NICT} P_1 + I_{ICT} P_2 = \bar{H} W_H + K_{NICT} R_{NICT} + K_{ICT} R_{ICT}, \tag{19}$$

[8] Note that with $\theta = \rho = 0$ we obtain the familiar Cobb-Douglas preferences with $V = \log\left(\frac{e_j}{P_1^{\nu} P_2^{1-\nu}}\right)$, while $\nu = 0$ reduces the preferences to a one good economy with CRRA preferences. Boppart (2014) shows that $e_j^{\rho} \geq \frac{1-\rho}{1-\theta} \nu P_1^{\theta} P_2^{\rho-\theta}$ ensures that preferences are valid.

where R_{ICT}, R_{NICT} are interest rates and e_H is the total expenditure of the household.

We apply Roy's identity to the indirect utility function and obtain the Marshallian demand functions

$$C_{1,H} = \nu \frac{e_H}{P_1} \left(\frac{P_2}{e_H}\right)^{\rho} \left(\frac{P_1}{P_2}\right)^{\theta}, \qquad C_{2,H} = \frac{e_H}{P_2} \left(1 - \nu \left(\frac{P_2}{e_H}\right)^{\rho} \left(\frac{P_1}{P_2}\right)^{\theta}\right) \tag{20}$$

Additionally, the high-skill household solves an intertemporal optimization problem, choosing expenditures and investment. Denoting by $\beta \in [0, 1]$ the discount factor, the value function can be expressed as

$$W(K_{NICT}, K_{ICT}) = \max V(P_1, P_2, e_H) + \beta(1 + g_H)\mathbb{E}\left[W(K'_{NICT}, K'_{ICT})\right], \tag{21}$$

where g_H is the exogenous growth rate of the H-household as specified above. Optimization results in the Euler equations

$$\begin{aligned} \left(\frac{P'_2}{P_2}\right)^{\rho} \left(\frac{e'_H}{e_H}\right)^{1-\rho} &= \beta\,(1 + g_H)\,\frac{(P'_2(1 - \delta'_{ICT})\mu'_{ICT} + R'_{ICT})}{\mu_{ICT} P_2}, \\ \left(\frac{P'_2}{P_2}\right)^{\rho} \left(\frac{e'_H}{e_H}\right)^{1-\rho} &= \beta\,(1 + g_H)\,\frac{(P'_1(1 - \delta'_{NICT})\mu'_{NICT} + R'_{NICT})}{\mu_{NICT} P_1} \end{aligned} \tag{22}$$

which we can also write jointly as a no-arbitrage condition

$$\frac{P'_2}{P_2}\left(\frac{\frac{R'_{ICT}}{P'_2} + \mu'_{ICT}(1 - \delta'_{ICT})}{\mu_{ICT}}\right) = \frac{P'_1}{P_1}\left(\frac{\frac{R'_{NICT}}{P'_1} + \mu'_{NICT}(1 - \delta_{NICT})}{\mu_{NICT}}\right)$$

This condition states that effective returns on ICT capital have to equal effective returns on non-ICT capital.

Low-skill workers are excluded from capital markets and are hand-to-mouth consumers. Each worker j maximizes eq. (18) subject to

$$\underbrace{C_{1,j}P_1 + C_{2,j}P_2}_{e_j} = \max\{\lambda_{R,j}W_R, \lambda_{M,j}W_M.\} \tag{23}$$

Individual-level Marshallian demands can be expressed correspondingly to eq. (20) and are aggregated to

$$C_{1,o} = \nu \frac{e_o}{P_1} \left(\frac{P_2}{e_o}\right)^{\rho} \left(\frac{P_1}{P_2}\right)^{\theta} \Delta_o, \qquad C_{2,o} = \frac{e_o}{P_2} \left(1 - \nu \left(\frac{P_2}{e_o}\right)^{\rho} \left(\frac{P_1}{P_2}\right)^{\theta}\right) \Delta_o \tag{24}$$

where $o = R, M$, such that e_o is total expenditure of R or M workers, respectively. $\Delta_o = \int_0^o \left[e_j / e_o\right]^{\rho} dj$ is a measure of dispersion of consumption expenditures across workers j within occupation o (see Boppart, 2014).

Boppart (2014) also shows that the elasticity of substitution between goods 1 and 2, σ_j, is worker-

specific and given by

$$\sigma_j = 1 - \theta - \frac{\nu \left(\frac{P_1}{P_2}\right)^{\theta}}{\left(\frac{e_j}{P_2}\right)^{\rho} - \nu \left(\frac{P_1}{P_2}\right)^{\theta}} (\theta - \rho).$$

3.4 Market Clearing

The two capital markets clear

$$K_{NICT} = K_{NICT,1} + K_{NICT,2}, \qquad K_{ICT} = K_{ICT,1} + K_{ICT,2}. \tag{25}$$

The markets for consumption goods clear

$$C_1 = C_{1,H} + C_{1,R} + C_{1,M}, \qquad C_2 = C_{2,H} + C_{2,R} + C_{2,M}. \tag{26}$$

and finally the labor markets clear

$$\bar{H} = H_1 + H_2, \qquad \bar{R} = R_1 + R_2, \qquad \bar{M} = M_1 + M_2 \tag{27}$$

3.5 Equilibrium

Given paths for μ_{NICT}, μ_{ICT} g_{A1}, g_{A2}, the distribution of low-skill abilities $\Gamma(\lambda_R, \lambda_M)$ and the exogenous supply of high- and low-skill workers $\bar{H}$, $\bar{L}$, the equilibrium is defined as a set of wages W_H, W_R, W_M, prices P_1, P_2 and returns R_{ICT}, R_{NICT} such that the first order conditions of the firms, eqs. (7), (8), (9), (10), (11), are satisfied, the first order conditions of the workers, eqs. (20), (22) and (24) are satisfied and markets clear, eqs. (25), (26) and (27).

Initially, we set $P_2 = 1$ without loss of generality. Subsequent nominal values are normalized such that inflation for the median household is zero. This implies deflating with a Törnqvist price index where the weights refer to the expenditure on goods 1 and 2 for this household.

4 Calibration and Simulation

Mirroring the empirical analysis, we consider changes in the model economy between 1960 and 2017. We assume that the economy is at a pre-digitalization steady state in 1960, and converges to a new steady state after 2017. Between 1960 and 2017, three kinds of exogenous processes take place: changes in the the rate of transformation of output into capital in both sectors, changes in sector-level TFP, and changes in the supply of high- and low-skill labor. Our calibration builds on the assumption that these exogenous processes fully explain changes in the demand and supply of factor inputs and consumption goods and corresponding equilibrium prices. This allows us on the one hand to study income and price changes in their entirety and assess their relative contributions to welfare changes. On the other hand, we can assess the importance of an individual channel

(most notably digitalization) by studying its effect in isolation.

We first calibrate all parameters that can be identified directly from the data and set some others in line with the literature. All remaining parameters are internally calibrated by matching empirical moments that are informative for the corresponding parameters, following the Simulated Method of Moments (McFadden, 1989). We show that our calibrated parameters are close to existing estimates in the literature.

4.1 Externally Chosen Parameters

Sector Classification

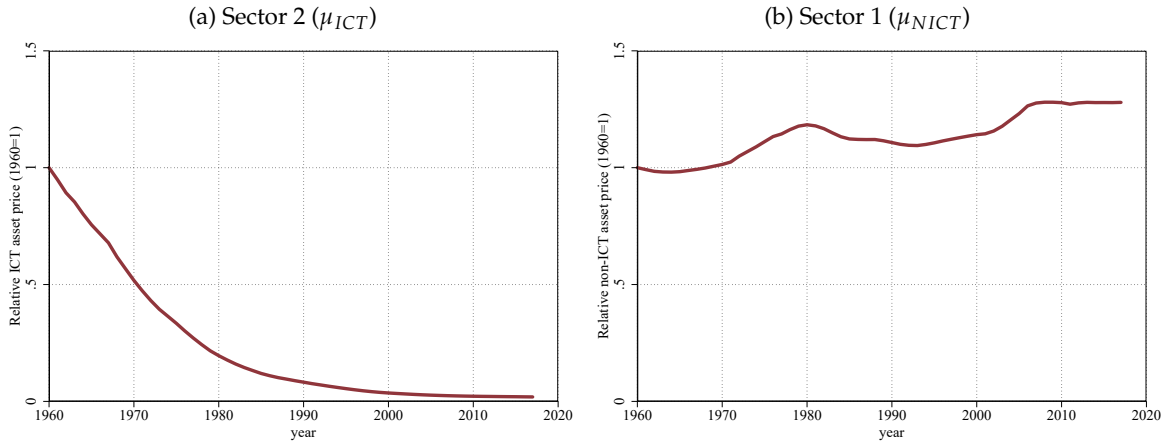
We sort the product categories from Section 2 into an ICT and a non-ICT composite good by clustering their ICT shares in 1996-98, which lies approximately in the middle of the sample. Product categories with an ICT share above a cut-off value of 0.065 get classified as ICT-intensive. This set comprises 38 product categories, which correspond to seven BEA industries. As Table 2 shows, ICT-intensive industries are first and foremost industries relating to computer hardware and software production and distribution. However, financial services and management activities are also on the list. The value-added share of this group of industries is 7.6% in 1960 and increases to 13.2% in 2017.

Our approach classifies industries based on the inputs used in the production process. This corresponds closely to our definition in Section 2 and is therefore our preferred classification. An alternative approach, equally in line with eq. (12), would be to classify them according to whether they produce ICT capital. We address this in Appendix B.1. Industries defined as non-ICT in our input-based sorting cover 80% of the production of non-ICT assets while all ICT capital is produced by the ICT sector. This implies that the two classification methods result in similar allocations. In Appendix B.1, we also compare our classification to the classification of manufacturing, high- and low-skill services by Buera and Kaboski (2012).

Table 2: ICT-intensive industries

BEA industry	industry name	ICT share 1996-98
3340	Computer and electronic products manufacturing	16.3%
5415	Computer systems design and related services	11.9%
5110	Publishing industries	10.1%
5210, 5220	Federal Reserve banks, credit intermediation, and related activities	8.4%
5320	Rental and leasing services and lessors of intangible assets	8.4%
5140	Data processing, internet publishing, and other information services	7.2%
5500	Management	6.7%

Figure 8: Rate of technological progress



Note: μ_{ICT} is the price of ICT capital divided by the price of the ICT-intensive good. μ_{NICT} is the price of non-ICT capital divided by the price of the non-ICT-intensive good. *Source:* NIPA, BEA and own calculations.

Advances in Digital Technology

Progress in ICT technology works through a decline in μ_{ICT}, the rate of transformation of output into ICT capital. To capture this, we take the price of ICT capital, p_{ICT}, from Section 2.1 and divide it by the price of output. Figure 8 (a) shows a strong decline in μ_{ICT}. We feed a smoothed version of this series into our simulation as an exogenous change in ICT technology. The rate of transformation for non-ICT capital, μ_{NICT}, is derived accordingly. As shown in Figure 8 (b), it increased by about 25%.

Skill Supply and Occupational Choice

We feed in the supply of high-skill and low-skill labor, $\bar{H}$ and $\bar{L}$ as exogenous processes, using data from the Decennial Census and the American Community Survey (see Appendix A.2 for details). We define high-skill workers as workers in non-routine cognitive occupations (following the occupational classification of Autor and Dorn (2013)) and low-skill workers as all others. During our sample period,the relative supply of high-skill labor increases substantially from a share of 23.7% in 1960 to 40.2% in 2017. Appendix Figure A2 (a) shows the annual growth rates.

While high-skill workers always work in high-skill professions, low-skill workers decide whether to work in manual or routine occupations based on their individual abilities and the equilibrium wage (eq. (15)). They can switch occupation and sector at any time and without cost. We specify the distribution of routine and manual abilities, $\Gamma(\lambda_R, \lambda_M)$ to be jointly log-normal, $\Gamma(\lambda_R, \lambda_M) \sim \text{Lognormal}\left(\begin{bmatrix}\mu_R \\ \mu_M\end{bmatrix}, \begin{bmatrix}\sigma_R^2 & \rho_{R,M} \\ \rho_{R,M} & \sigma_M^2\end{bmatrix}\right)$. There are five parameters to specify: two standard deviations, two means and one correlation. Following Jaimovich et al. (2021), we normalize the means to 1 and assume that the correlation between the two ability draws is zero [9].

[9]Normalizing the means to 1 is inconsequential, as they are essentially just a scaling factor over simulated wages. The correlation between the two abilities, $\rho_{R,M}$ cannot be identified in the data. Jaimovich et al. (2021) experimented with different values and concluded that the choice is inconsequential, therefore opting for $\rho_{R,M} = 0$ for convenience.

We compute the standard deviations of wages, σ_M and σ_R from the ACS and Census data. The data only contains wage information for the occupation actually chosen by the individuals, which means that our sample is truncated. We obtain the standard deviations of the full-distribution of routine and non-routine manual wages via a truncated maximum likelihood estimation. Appendix Figure A2 (b) shows the resulting numbers.

We require the ratio between routine and manual wages to be such that the model exactly produces the routine worker share observed in the data. Concretely, we make a guess for the routine-manual wage premium, draw individual productivities using the above specified process, and let individuals decide on whether to work in R- or M-jobs using eq.s (16) and (17). We then compare the resulting supply of routine and manual workers with the actual supplies from the data, and iterate on our guess for the routine-manual wage premium until we exactly replicate the observed worker distribution. We use the resulting manual-routine wage ratio as a data moment to target (see below). Note that individual wages in the data correspond to the model's $\lambda_{j,R}W_R$ and $\lambda_{j,M}W_M$, thus including the ability parameters. We calibrate the income distribution parameters Δ_o that are used for the Marshallian demands in eq. (24) by creating a quantile function from $\Gamma'(\lambda_R, \lambda_M)$. This function then represents the distribution of wages across the skill distribution.

Production Function

All parameters of the production function are assumed constant over time. We use values for the 1996-98 period where available. The depreciation rates δ_{NICT}, δ_{ICT} are taken from the BEA. The non-ICT capital shares α_i are computed as one minus the ICT share minus the labor share in each sector i. The share of manual work ψ_i is calculated by multiplying the share of non-routine manual workers from the 2000 Census with the labor share from Elsby et al. (2013).

4.2 Internally Calibrated Parameters

The remaining parameters of the production function and the preference parameters are calibrated via the Simulated Method of Moments. Appendix B.2 provides further details on the calibration procedure and the solution algorithm. To keep the number of parameters and moments tractable, we make the simplifying assumption that both sectors have the same elasticities of substitution, $\eta_1 = \eta_2$ and $\epsilon_1 = \epsilon_2$[10]. This leaves us with 11 parameters to be calibrated. Next to the two elasticities of substitution, these are the four weights of the production function (γ_i, ϕ_i), the parameters in the indirect utility function ν, ρ and θ, as well as the annual TFP growth rate g_{Ai} in each sector.

We aim to match 11 moments, such that the system is exactly identified under the assumption of moments orthogonality. Most moments are expressed in percentage changes between 1960 and 2017. As explained in more detail in Appendix B.3, each of the moments is particularly informative about one or more parameters. For calibrating the parameters in the production function,

[10]This means that both sectors have the same capability to substitute among input factors but may still have different weights γ_i, ϕ_i. We see in figure 3 that all sectors experience a growth in their ICT share, and sectors that have been the most ICT intensive in 1960 are still the most ICT intensive now. This motivates us to choose the same elasticities ϵ and η across sectors. See Klump and de La Grandville (2000) for a more general discussion on how weights and elasticities are related.

we choose the following moments: First, two moments about relative hourly wages, which are informative about the relative use of the different types of labor and are particularly closely related to the elasticities of substitution in the production functions. These are the skill premium (wage premium of high-skill relative to routine work) and the routine-manual premium (wage premium of routine relative to manual work). We obtain the skill premium directly from ACS and Census data, and use the routine-manual wage premium from a calibration based on the same data sources as described in Section 4.1. Additionally, we use the labor share in value-added, the rate of return to ICT capital and the ICT intensity as moments for the production function, which are informative about the use of labor vs. capital inputs. The labor share is taken from Elsby et al. (2013), to which we apply HP filtering and extrapolate the series between 2013-2017. The rate of return to ICT capital and the ICT intensity are directly taken from Section 2. We also choose the price of good 1 relative to good 2, compiled as described in Section 4.1, which is relevant both for the sectoral differences in the production functions and the parameters of the utility function.

For calibrating the parameters in the utility function, we choose the difference in the expenditure share of good 2 between the 99th and the 1st percentile in the income distribution, the aggregate quantities consumed of good 1 relative to good 2, and the expenditure share of good 2 for the median household. All of these are based on CEX and NIPA data as described in Section 2. Finally, to match the sector growth rates of TFP, we choose the (real) GDP per capita and the value-added share of sector 2, both taken from the BEA.

While each moment is particularly important for one set of parameters, there is not a one-to-one mapping, and changes in one parameter affect several outcomes at once. Appendix B.2 provides further details on the calibration procedure, the solution algorithm and how the moments link to the parameters.

4.3 Calibration Outcomes

Table 3 compares the targeted data moments with the model. Most moments are well matched within a 10% range. Importantly, we match the relative price decline and the difference in expenditure shares very well, which is crucial to gauge the price effect. Table 4 summarizes the calibrated parameters. The elasticity of substitution in the inner nest (η) is 1.29 and suggests that routine labor and ICT capital are substitutes. Our parameter value lies close to the value of Vom Lehn (2020), who provides an estimate of 1.35. Other papers, such as Eden and Gaggl (2018), Bárány and Siegel (2021) or Jaimovich et al. (2021) have estimated elasticities closer to 2.

The routine-ICT composite and high-skilled labor emerge as complementary to each other through the low value of $\epsilon = 0.80$. Again, our calibrated value lies close to estimates used in other papers: Vom Lehn (2020) uses a value of 0.83. In a slightly different set-up of the production function, Bárány and Siegel (2021) assume a value of 0.6 for the elasticity of substitution between the ICT-routine aggregate and two types of labor. Jaimovich et al. (2021) use a value of 0.32.

We see not surprisingly that $\phi_1 > \phi_2$, as sector 2 is more ICT-intensive whereas sector 1 relies more on routine labor input. In the middle nest, the input generated using these two production factors is more important in sector 2, whereas sector 1 places a high weight on high-skill labor. As sector

Table 3: Calibration fit

Data moment	Source	Data	Model
Skill premium, 1960-2017	ACS	39.9	37.2
Relative price, 1960-2017	NIPA	244.2	237.6
ICT intensity, 1960-2017	see Sec. 2	643.1	596.6
Relative consumption, 1960-2017	NIPA	-85.4	-89.2
Labor share, 1960-2017	EHS (2013)	-11.7	-7.2
Rate of return to ICT capital, 1960-2017	see Sec. 2	-97.8	-98.1
Median increase in exp. share good 2, 1960-2017	see Sec. 2	70.7	65.6
Percent diff., exp. share good 2 Top 1 Bot 1, 2017	see Sec. 2	23.3	24.8
Routine-manual wage premium, 1960-2017	see Sec. 4.1	-31.7	-31.6
GDP per capita, 1960-2017	BEA	281.6	286.3
Value added share sector 2, 1960-2017	BEA	71.4	60.1

Note: All moments are expressed in percentage changes, except for the difference in expenditure shares. EHS refers to Elsby et al. (2013).

1 relies more on non-ICT capital ($\alpha_1 > \alpha_2$) and on manual work ($\psi_1 > \psi_2$), high-skill labor plays a less important role than the large value of γ_1 suggests.

A positive value of ρ indicates that the ICT good is a luxury good, in line with the empirical evidence. $\theta < 0$ implies that when the relative price of a good increases, the share of expenditure spent on this good decreases. This means that the two goods are substitutes. Indeed, the elasticity of substitution between ICT and non-ICT goods for the average worker is $\sigma = 1.57$. With an elasticity of substitution above one, the sector that experiences a relative price increase (i.e. the non-ICT sector) shrinks in terms of expenditure shares.[11] To our knowledge, our paper is the first to provide estimates of the elasticity of substitution between ICT and non-ICT goods.

On the TFP parameters, the best fit is achieved when the ICT intensive sector has a positive TFP growth rate of 0.12% per year, whereas the low ICT sector has a TFP growth rate that is positive but close to zero. These numbers are slightly lower than numbers reported in Inklaar et al. (2005). They find that overall TFP growth in the US was 0.26% between 1979-95 and 1.05% in 1995-2000. They also show that most of this growth stems from ICT industries, whereas some non-ICT industries contributed negatively to TFP growth. Jorgenson and Stiroh (1999) report aggregate TFP growth rates of 0.34% for 1973-90 and 0.23% for 1990-96.

We show the validity of the estimated model by contrasting some non-targeted model moments with data moments in Appendix B.4. There we also verify that the overall distribution of income, price and expenditure changes matches the data well along the income distribution.

[11] Compare this to Boppart (2014) who considers a goods-services dichotomy and finds these two to be complements.

Table 4: Calibrated parameters

Symbol	Value	Description	Source
Production function: Inner nest			
ϕ_1	0.756	Weight of R in 1	calibrated
ϕ_2	0.523	Weight of R in 2	calibrated
$\eta_1 = \eta_2$	1.288	El. of Subst between R and ICT	calibrated
Production function: Middle nest 1			
γ_1	0.662	Weight of H in sector 1	calibrated
γ_2	0.319	Weight of H in sector 2	calibrated
$\epsilon_1 = \epsilon_2$	0.802	El. of Subst betw H and AW	calibrated
Production function: Middle nest 2			
ψ_1	0.129	Income share of M in sector 1	ACS and EHS
ψ_2	0.029	Income share of M in sector 2	ACS and EHS
Production function: Final Production			
α_1	0.372	Capital share in sector 1	BEA
α_2	0.200	Capital share in sector 2	BEA
δ_{ICT}	0.14	Depreciation rate ICT capital	BEA
δ_{NICT}	0.06	Depreciation rate non-ICT capital	BEA
Preferences			
ν	0.517	Expenditure share parameter good 1	calibrated
θ	-0.306	Substitution parameter	calibrated
ρ	0.043	Income elasticity parameter	calibrated
β	0.97	Discount factor	
Growth in skill supply			
g_H	0.027	mean growth rate of H supply	ACS
g_L	0.013	mean growth rate of L supply	ACS
TFP growth			
g_{A1}	0.0006	TFP growth in sector 1	calibrated
g_{A2}	0.0012	TFP growth in sector 2	calibrated
Ability distribution			
σ_M	0.625	mean std. dev. of M ability	estimated
σ_R	0.534	mean std. dev. of R ability	estimated
μ_M	1	mean M ability	normalization; JSSL
μ_R	1	mean R ability	normalization; JSSL
$\rho_{R,M}$	0	correlation M, R abilities	JSSL
Dispersion of consumption expenditures across workers			
Δ_M	1.296	Consumption dispersion, M	estimated
Δ_R	1.297	Consumption dispersion, R	estimated

Note: EHS refers to Elsby et al. (2013). JSSL refers to Jaimovich et al. (2021). Standard deviations of ability distributions and growth rates of labor supply are time-varying and the table shows averages; the full time series can be found in Appendix Figure A2.

5 Results

In this section, we present simulation results for the effect of digitalization on prices, income and consumption of workers. We show that price changes vary substantially across the income distribution and matter for consumption inequality. To estimate the importance of the price effect for welfare, we carry out an equivalent variation that assigns a monetary value to the price and income changes in each income group. The overall welfare effects of digitalization emerge as a combination of a U-shaped pattern of income polarization, combined with relative price changes that are strictly downward-sloping in income. As a result, welfare changes are J-shaped along the income distribution. Lastly, two counterfactual exercises highlight the importance of digitalization and non-homothetic preferences for welfare effects.

5.1 Welfare Changes

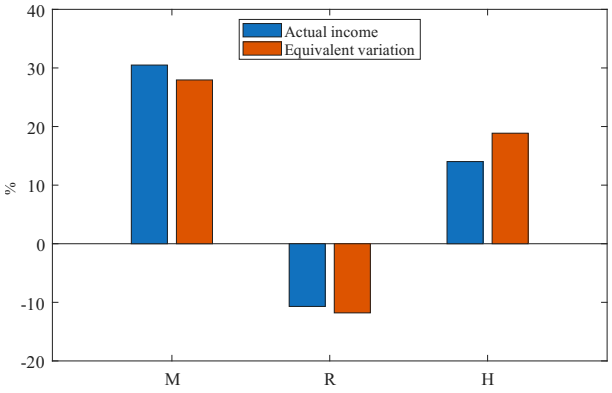
To assess welfare changes along the income distribution, we use the estimated utility function together with the movement in key variables to calculate an equivalent variation for each income percentile. An equivalent variation asks how much additional income workers need to be given in 1960 – given the initial prices – to receive the same utility as they are getting with income and prices in 2017. If the equivalent variation is lower than the actual income increase, relative price changes have been disadvantageous. Using the model, we can not only calculate total welfare changes along the income distribution, but also decompose these into a price and an income effect, and thus assess how important price changes are relative to income changes.

Figure 9 (a) compares the equivalent income with the actual income per worker for manual workers (W_M), routine workers (W_R) and high-skill workers ($W_H + (K_{NICT}R_{NICT} + I_{ICT}R_{ICT})/\bar{H}$). High-skill workers also get income from their ownership of capital, however, this does not make a big difference for relative income gains. We discuss the role of capital income in more detail in Appendix Section B.5. Actual and equivalent incomes are expressed as deviation from mean income growth and deflated with the price index for the median household, as discussed in Section 3.5. Figure 9 (b) highlights the polarizing effect of income changes for the different percentiles of the low-skill group. Here, we simulate 100,000 low-skill workers by drawing manual and routine ability pairs λ_R, λ_M for each and sorting them into the occupation that pays the highest wage. We then rank them according to their income. Owing to the changing ability distribution over time, we draw separate samples for 1960 and 2017. The heterogeneous group of low-skill workers corresponds to income percentiles 1 to 77 in our model, which is their share in the 1960 data. Percentiles 78 to 100 are high-skill workers. Given that individual wages are a combination of equilibrium wages and individual ability draws, each percentile will include a mix of routine and manual workers. Therefore, the differences across the low-skill distribution in panel (b) are much less pronounced than those between groups in panel (a). The blue dots show how much more income workers actually receive in the model in 2017 relative to the mean increase, whereas the red bars show the equivalent variation.

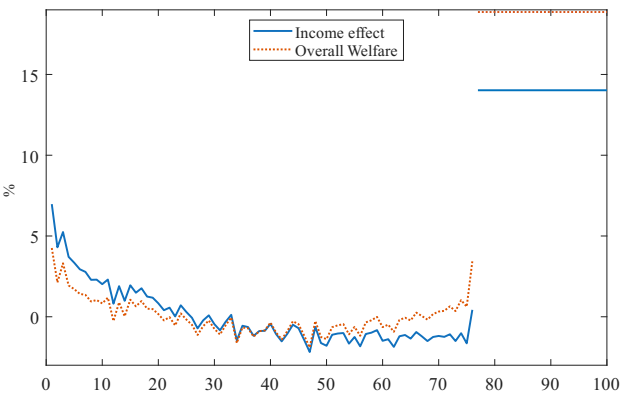
Considering equivalent income changes, we see that both manual and routine workers experience disadvantageous price changes, while high-skill workers now emerge as the clear winners, being

Figure 9: Income changes and equivalent variation, deviation from mean income growth

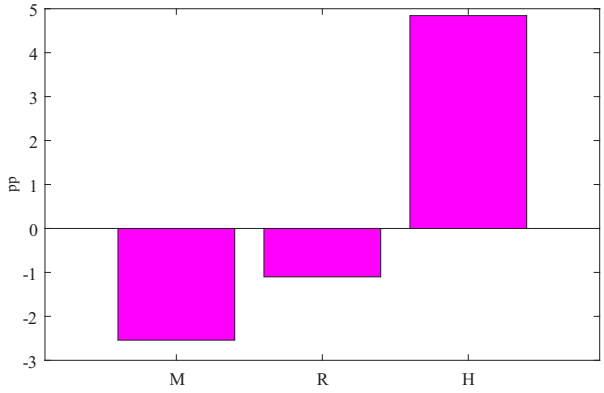
(a) by group



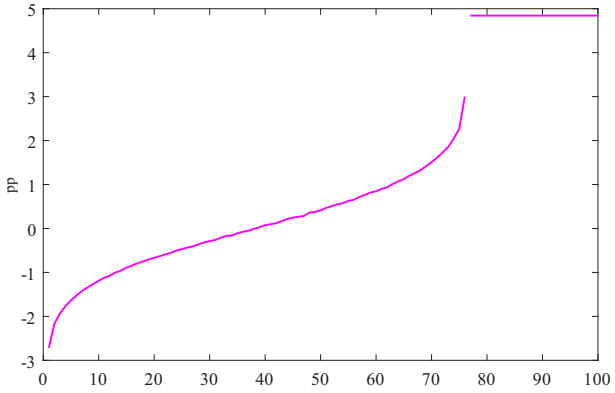
(b) by percentile



(c) EV - income by group



(d) EV - income by percentile



Note: Panel (a) shows deviations from the average income increase of actual income in 2017 (blue) and equivalent variation (red) for the manual, routine and high-skill worker groups. Panel (b) shows deviations from the average income increase of actual income (blue line) and the equivalent variation (red dotted line) for percentiles of the income distribution of low-skill workers. Panel (c) and (d) shows the difference between the equivalent variation and actual income.

the primary consumers of luxury goods. When considering income percentiles in panel (b), the large income gains at the very bottom get diminished by a loss of purchasing power, while at the upper part of the low-skill distribution, relative income losses get mitigated by lower ICT prices. Overall, the equivalent variation curve is substantially flatter than the income curve, and in terms of welfare the very bottom and top of the low-skill group are much closer than the pure income effect suggests. Taken together, income and price changes imply that welfare increases more at the top than at the bottom. The overall welfare change resembles more a J-shape than a U-shape.

How much of the relative welfare effect is driven by changes in relative prices? The equivalent variation (EV) captures both the income effect and the price effect. The contribution of the price effect can be obtained as the difference between the equivalent variation (red) and the pure increase in expenditures (blue) (Baqaee and Burstein, 2023). Figures 9 (c) and (d) quantify the role of the price channel. Panel (c) reports differences across groups in percentage points. Between 1960 and 2017, changes in relative purchasing power translate into a decline in welfare-equivalent income of 2.5 percentage points for manual and 1 percentage points for routine workers, whereas

high-skill workers experience an increase of around 5 percentage points. Panel (d) considers the whole distribution. Above-average income gains of 7% for the bottom group are reduced by 2.75 percentage points to welfare-equivalent income gains of 4.25%. This implies that the price channel reduces the relative welfare increase by roughly 30% for that group. At the top of the distribution, by contrast, the price effect amplifies the relative income gain by 5 percentage points from 14% to 18.9%, corresponding to an additional welfare increase of about 35%.

A few remarks are in order. In our model, price changes translate into welfare changes. However, this does not necessarily have to be the case, but depends crucially on the specification of preferences and the product space. Oberfield (2023) considers an alternative set-up where higher-income households experience lower price changes, but where these do not translate into higher welfare gains for this group. The reason for the difference is that in his setting, products have a life-cycle: Any product is initially a luxury and then becomes a necessity as aggregate income grows. Therefore, low-income households not only benefit from productivity increases in current necessities, but even more so, they benefit from productivity increases in current luxuries which they get to consume later in the product's life-cycle. Commonly used inflation measures only capture the current rate of change in prices, and in Oberfield's setting are therefore misleading about welfare differences across income groups. In our setup, all households consume all goods, and at any point in time, the share of luxuries in the consumption basket of rich households is larger than the share in the basket of poor households. As innovation happens in luxuries, richer households are better off at any time. The two frameworks can be reconciled if one considers Oberfield's set-up as a description of what happens within each of the two sectors of our model, e.g. a first-generation mobile phone being replaced by a state-of-the-art smartphone as consumers get richer. Still, everyone consumes mobile phones, and richer households spend a larger share of their income on mobile phones at any time.

5.2 Digitalization and Welfare

To what extent are welfare changes caused by digitalization? Would the income and price effect look different in the absence of digitalization? To answer these questions, we carry out a counterfactual analysis, where we assume that digitalization had not taken place. Concretely, we set $\mu_{ICT} = 1$ at all times. The rest of the set-up is kept unchanged.

Table 5 shows the new set of generated moments in this counterfactual scenario. Most directly affected by digitalization is the return to ICT capital, which does not change anymore. The higher user costs of ICT capital relative to the baseline simulation result in a different organization of production. The ICT intensity increases only slightly, and the value added share of sector 2 stays almost the same. As a result, the price of good 1 relative to good 2 barely moves. We also see a strong effect on GDP, which grows substantially less now. This shows that the productivity gain from digitalization was an important driver of economic growth in the baseline model. However, economic growth in the model hinges crucially on the interaction of all three exogenous processes (changes in the productivity of capital, TFP growth and changes in labor supply) and is therefore not additive. In Appendix B.6 we show that muting the labor supply increase also leads to a large decline in GDP per-capita growth. The effect of digitalization on GDP growth mirrors e.g.

Table 5: Counterfactual no digitalization ($\mu_{ICT} = 1$) model vs. data

Moment	Data	Model
Skill premium, 1960-2017	39.9	-64.9
Relative price, 1960-2017	244.2	4.1
ICT intensity	643.1	38.4
Relative consumption, 1960-2017	-85.4	-7.3
Labor share, 1960-2017	-11.7	-3.9
Rate of return to ICT capital, 1960-2017	-97.8	0.0
Median increase in exp. share good 2, 1960-2017	70.7	1.8
Percent diff., exp. share good 2 Top 1 Bot 1, 2017	23.3	70.5
Routine-manual wage premium, 1960-2017	-31.7	61.3
GDP per capita, 1960-2017	281.6	12.7
Value-added share of sector 2, 1960-2017	71.4	3.8

Note: All moments are expressed in percentage changes, except for the difference in expenditure shares.

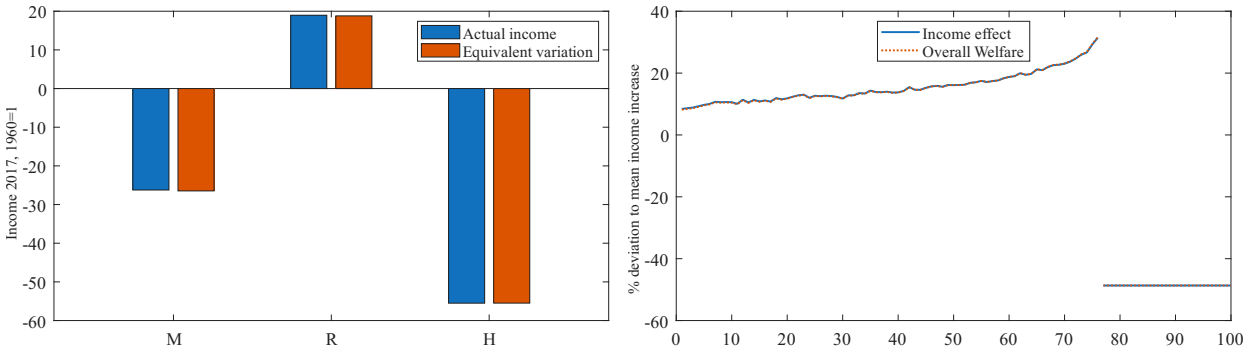
Jorgenson and Stiroh (1999), who argue that computers played an important role for economic growth in the 1990s, whereas TFP growth was depressed.

Changes in the production structure also affect labor demand, and thereby the wages of households. The skill premium, instead of increasing, now decreases over time, whereas the premium of routine over manual wages increases. The literature has postulated that over the last decades, the evolution of the skill premium was due to a combination of technology-driven increases in the relative demand for high-skill workers and an increasing college attendance rate, thus increasing the supply of high-skill cognitive workers (see Goldin and Katz (2008) for an overview). When muting the labor demand channel, but keeping the labor supply channel intact, the skill premium therefore decreases. The opposite is the case if we mute the labor supply channel (see Appendix Section B.6). The change in wage premia manifests in large differences in income changes, as shown by the blue bars and dots in Figure 10. Now those best off in terms of income are the routine workers, whereas the high-skill workers actually see a decrease in their labor income in 2017 compared to 1960. Thus, polarization is a result of digitalization.

The share of expenditures spent on good 2 is now growing much less as there is less substitution. As the red bars and dots in Figure 10 show, in a world without digitalization there is no meaningful price effect, the equivalent variation tracks income changes. Thus, the price effect in absence of digitalization is rather flat, and welfare changes largely mirror income changes.

Figure 10: Income changes and equivalent variation: $\mu = 1$

(a) by group (b) by percentile



Note: This figure replicates Figure 9 (a) and (b) for the no-digitalization counterfactual. Panel (a) shows actual income in 2017 (blue) and equivalent variation (red) for the manual, routine and high-skill worker groups. Panel (b) shows actual income increase and the equivalent variation for percentiles of the income distribution of low-skill (manual and routine) workers. Income values are relative to 1960.

5.3 The Role of Non-Homothetic Preferences

We have highlighted throughout the paper that non-homothecity is crucial for consumption price inequality and therefore the existence of the price channel. If households had the same consumption patterns, there would be no divergence of price indices. Our model allows us to gauge the importance of this channel by muting non-homothetic preferences, setting the parameter ρ to zero. All other parameters and shocks remain unchanged. It therefore creates a counterfactual world in which expenditure shares on ICT consumption are the same along the income distribution.

Table 6 summarizes how the matched moments change. By eliminating non-homotheticity in preferences, the differences in expenditure shares disappear. This implies that the ICT good is no longer a luxury good. As a result, the increase in the aggregate expenditure share is lower and exclusively driven by substitution due to lower digital prices. Furthermore, the increase in value-added share of the ICT sector is also lower. ICT intensity – the share of ICT capital in the total

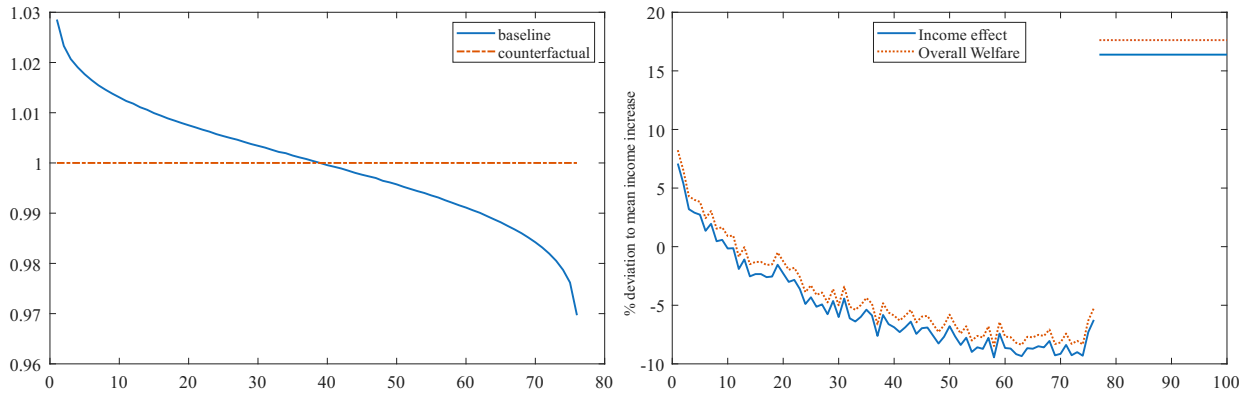
Table 6: Counterfactual homothetic preferences ($\rho = 0$) model vs. data

Moment	Data	Model
Skill premium, 1960-2017	39.9	56.1
Relative price, 1960-2017	244.2	208.2
ICT intensity	643.1	317.1
Relative consumption, 1960-2017	-85.4	-86.7
Labor share, 1960-2017	-11.7	-5.3
Rate of return to ICT capital, 1960-2017	-97.8	-98.1
Median increase in exp. share good 2, 1960-2017	70.7	28.0
Percent diff., exp. share good 2 Top 1 Bot 1, 2017	23.3	0.0
Routine-manual wage premium, 1960-2017	-31.7	-40.9
GDP per capita, 1960-2017	281.6	391.6
Value-added share of sector 2, 1960-2017	71.4	25.5

Note: All moments are expressed in percentage changes, except for the difference in expenditure shares.

Figure 11: Equivalent variation and prices

(a) Price index with $\rho = 0$ (b) EV with $\rho = 0$



Note: The left panels show the price index for each income percentile with homothetic preferences in the baseline (solid blue) and the counterfactual (dashed red). The right panel shows the equivalent variation of the baseline (blue circles), the equivalent variation of the counterfactual (red dots) for percentiles of the income distribution of low-skill (manual and routine) workers.

capital stock – goes down substantially, which also coincides with a less pronounced decrease of relative prices. While the main changes happen for the price effect, income polarization gets also more pronounced due to changes in the demand for the different types of capital and goods.

Prices no longer affect consumers differently when expenditure shares are the same, as indicated in Figure 11 (a). As a consequence, the welfare effects in Figure 11 (b) now fully mirror the income effects. Poor households are not adversely affected by the price effect in the counterfactual. At the same time, the stronger U-shaped income polarization hits the middle of the income distribution especially hard and makes their welfare increase lower in the counterfactual.

We conclude that the inclusion of non-homothetic preferences together with the relative price movements are crucial to understand the relation between income and overall welfare effects. The divergence of price indices that we observe in the data can indeed have meaningful implications for welfare.

6 Conclusion

This paper studies the impact of digitalization on prices and consumption along the income distribution. We provide an empirical estimate of the effect of digitalization on the price of consumption along the income distribution, and quantify price and welfare effects in a structural model. Our results show that the effects of digitalization on consumption inequality and welfare are not determined exclusively by income changes, but that relative price changes play an important role. While income changes benefit households both at the top and at the bottom of the income distribution, price changes are universally more favorable for higher-income households. In providing a first estimate of the size of this price channel, this paper offers a starting point for more follow-up research on the effect of digitalization on welfare and inequality, which could for example address the following questions: In how far does digitalization affect quality differences between relat-

ively similar products consumed by the poor or the rich? Has digitalization changed the nature of how people source and access goods? Furthermore, this paper has abstracted from explaining how advances in digital technology are generated. Incorporating learning by doing, the creation of new products or quality changes in existing products, the framework could be further refined.

References

Acemoglu, D. and P. Restrepo (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review 108*(6), 1488–1542.

Aghion, P., C. Antonin, X. Jaravel, and S. Bunel (2020). What are the labor and product market effects of automation? New evidence from France. *CEPR Discussion Paper 14443*.

Aguiar, M. and M. Bils (2015). Has consumption inequality mirrored income inequality? *American Economic Review 105*(9), 2725–2756.

Aguiar, M. and E. Hurst (2005). Consumption versus expenditure. *Journal of Political Economy 113*(5), 919–948.

Akerman, A., I. Gaarder, and M. Mogstad (2015). The skill complementarity of broadband internet. *Quarterly Journal of Economics 130*(4), 1781–1824.

Argente, D. and M. Lee (2021). Cost of living inequality during the great recession. *Journal of the European Economic Association 19*(2), 913–952.

Attanasio, O. P. and L. Pistaferri (2014). Consumption inequality over the last half century: some evidence using the new PSID consumption measure. *American Economic Review 104*(5), 122–26.

Attanasio, O. P. and L. Pistaferri (2016). Consumption inequality. *Journal of Economic Perspectives 30*(2), 3–28.

Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review 103*(5), 1553–1597.

Autor, D. H., L. F. Katz, and M. S. Kearney (2008). Trends in US wage inequality: Revising the revisionists. *Review of Economics and Statistics 90*(2), 300–323.

Baqaee, D. R. and A. Burstein (2023). Welfare and output with income effects and taste shocks. *Quarterly Journal of Economics 138*(2), 769–834.

Bárány, Z. L. and C. Siegel (2021). Engines of sectoral labor productivity growth. *Review of Economic Dynamics 39*, 304–343.

Barkai, S. (2020). Declining Labor and Capital Shares. *Journal of Finance 75*(5), 2421–2463.

BEA (2021). Updated digital economy estimates - June 2021. *BEA Report*.

Boppart, T. (2014). Structural change and the Kaldor facts in a growth model with relative price effects and non-Gorman preferences. *Econometrica 82*(6), 2167–2196.

Borusyak, K. and X. Jaravel (2021). The Distributional Effects of Trade: Theory and Evidence from the United States. Technical report, Available at SSRN, no. 3269579.

Brynjolfsson, E., A. Collis, W. E. Diewert, F. Eggers, and K. J. Fox (2019). GDP-B: accounting for the value of new and free goods in the digital economy. *NBER Working Paper No. 25695*.

Brynjolfsson, E., A. Collis, W. E. Diewert, F. Eggers, and K. J. Fox (2020). Measuring the impact of free goods on real household consumption. *AEA Papers and Proceedings 110*, 25–30.

Buera, F. J. and J. P. Kaboski (2012). The rise of the service economy. *American Economic Review 102*(6), 2540–2569.

Burstein, A., E. Morales, and J. Vogel (2019). Changes in between-group inequality: computers, occupations, and international trade. *American Economic Journal: Macroeconomics 11*(2), 348–400.

Byrne, D. and C. Corrado (2017a). Ict asset prices: Marshaling evidence into new measures. Technical report. FEDS Working Paper No. 2017-16.

Byrne, D. and C. Corrado (2017b). Ict prices and ict services: What do they tell us about productivity and technology? *International Productivity Monitor 33*, 150–181.

Cortes, G. M., N. Jaimovich, and H. E. Siu (2017). Disappearing routine jobs: Who, how, and why? *Journal of Monetary Economics 91*, 69–87.

Coyle, D. and D. Nguyen (2020). Free goods and economic welfare. *ESCoE Discussion Paper 2020-18*.

Eden, M. and P. Gaggl (2018). On the welfare implications of automation. *Review of Economic Dynamics 29*, 15–43.

Elsby, M. W. L., B. Hobijn, and A. Sahin (2013). The decline of the US labor share. *Brookings Papers on Economic Activity 2013*(2), 1–63.

Fajgelbaum, P. D. and A. K. Khandelwal (2016). Measuring the unequal gains from trade. *Quarterly Journal of Economics 131*(3), 1113–1180.

Fang, L., A. Hannusch, and P. Silos (2021). Luxuries, Necessities, and the Allocation of Time. *Working Paper*.

Gaggl, P. and G. C. Wright (2017). A short-run view of what computers do: Evidence from a UK tax incentive. *American Economic Journal: Applied Economics 9*(3), 262–94.

Goldin, C. and L. F. Katz (2008). *The race between education and technology*. Harvard University Press.

Goos, M. and A. Manning (2007). Lousy and lovely jobs: The rising polarization of work in Britain. *Review of Economics and Statistics 89*(1), 118–133.

Goos, M., A. Manning, and A. Salomons (2009). Job polarization in Europe. *American Economic Review 99*(2), 58–63.

Graetz, G. and G. Michaels (2018). Robots at work. *Review of Economics and Statistics 100*(5), 753–768.

Greenwood, J., Z. Hercowitz, and P. Krusell (1997). Long-Run Implications of Investment-Specific Technological Change. *American Economic Review 87*(3), 342–362.

Hémous, D. and M. Olsen (2022). The rise of the machines: Automation, horizontal innovation and income inequality. *American Economic Journal: Macroeconomics 14*(1), 179–223.

Herrendorf, B., R. Rogerson, and Á. Valentinyi (2013). Two perspectives on preferences and structural transformation. *American Economic Review 103*(7), 2752–2789.

Herrendorf, B., R. Rogerson, and Á. Valentinyi (2014). Growth and structural transformation. In *Handbook of Economic Growth*, Volume 2, pp. 855–941. Elsevier.

Horrowitz, K. J. and M. A. Planting (2009). Concepts and methods of the U.S. Input-Output Accounts. *BEA* (September 2006), 1–266.

Inklaar, R., M. O'Mahony, and M. Timmer (2005). ICT and Europe's productivity performance: industry-level growth account comparisons with the United States. *Review of Income and Wealth 51*(4), 505–536.

Jaimovich, N., I. Saporta-Eksten, H. E. Siu, and Y. Yedid-Levi (2021). The macroeconomics of automation: Data, theory, and policy analysis. *Journal of Monetary Economics 122*, 1–16.

Jaravel, X. (2019). The unequal gains from product innovations: Evidence from the US retail sector. *Quarterly Journal of Economics 134*(2), 715–783.

Jaravel, X. (2021). Inflation inequality: Measurement, causes, and policy implications. *Annual Review of Economics 13*, 599–629.

Jaravel, X. and D. Lashkari (2024). Measuring growth in consumer welfare with income-dependent preferences: Nonparametric methods and estimates for the united states. *The Quarterly Journal of Economics 139*(1), 477–532.

Jorgenson, D. W. and K. J. Stiroh (1999). Information technology and growth. *American Economic*

Review 89(2), 109–115.

Kaplan, G. and S. Schulhofer-Wohl (2017). Inflation at the household level. *Journal of Monetary Economics 91*, 19–38.

Karabarbounis, L. and B. Neiman (2014). The global decline of the labor share. *Quarterly Journal of Economics 129*, 61–103.

Karabarbounis, L. and B. Neiman (2019). Accounting for factorless income. *NBER Macroeconomics Annual 33*(1), 167–228.

Klump, R. and O. de La Grandville (2000). Economic growth and the elasticity of substitution: Two theorems and some suggestions. *American Economic Review 90*(1), 282–291.

Koh, D., R. Santaeulalia-Llopis, and Y. Zheng (2020). Labor share decline and intellectual property products capital. *Econometrica 88*(6), 2609–2628.

Krueger, D. and F. Perri (2006). Does income inequality lead to consumption inequality? Evidence and theory. *The Review of Economic Studies 73*(1), 163–193.

Kuhn, M. and J.-V. Ríos-Rull (2016). 2013 update on the US earnings, income, and wealth distributional facts: A view from macroeconomics. *Federal Reserve Bank of Minneapolis Quarterly Review 37*(1), 2–73.

Lashkari, D., A. Bauer, and J. Boussard (2024). Information technology and returns to scale. *American Economic Review 114*(6), 1769–1815.

Martinez, J. (2021). Putty-clay automation. *CEPR Discussion Paper 16022*.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica 57*(5), 995–1026.

Meyer, B. D. and J. X. Sullivan (2023). Consumption and income inequality in the United States since the 1960s. *Journal of Political Economy 131*(2), 247–284.

Michaels, G., A. Natraj, and J. van Reenen (2014). Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years. *The Review of Economics and Statistics 96*(1), 60–77.

Moll, B., L. Rachel, and P. Restrepo (2022). Uneven growth: Automation's impact on income and wealth inequality. *Econometrica 90*(6), 2645–2683.

NIPA (2023). NIPA Handbook: Concepts and methods of the U.S. National Income and Product Accounts.

Oberfield, E. (2023). Inequality and measured growth. Technical report. National Bureau of Economic Research.

Prettner, K. and H. Strulik (2020). Innovation, automation, and inequality: Policy challenges in the race against the machine. *Journal of Monetary Economics 116*, 249–265.

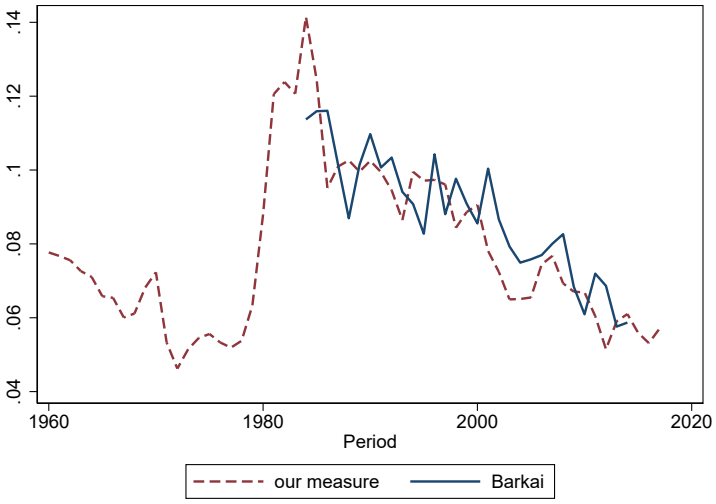
Veldkamp, L. and C. Chung (2019). Data and the aggregate economy. *Journal of Economic Literature*.

Vom Lehn, C. (2020). Labor market polarization, the decline of routine work, and technological change: A quantitative analysis. *Journal of Monetary Economics 110*, 62–80.

A Data Appendix

A.1 Additional Tables and Figures

Figure A1: Comparison of costs of capital to Barkai (2020)

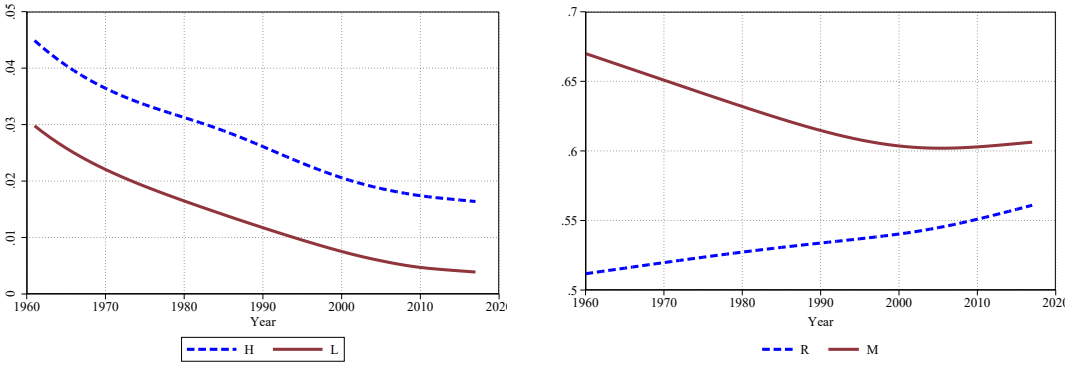


Note: The graph shows the costs of capital as calculated by $\tilde{r}_t$ in eq. (1) with the weighted average costs of capital of Barkai (2020), Figure 2A. *Source:* Barkai (2020), BEA and own calculations.

Figure A2: Labor supply and standard deviation of wages by worker type

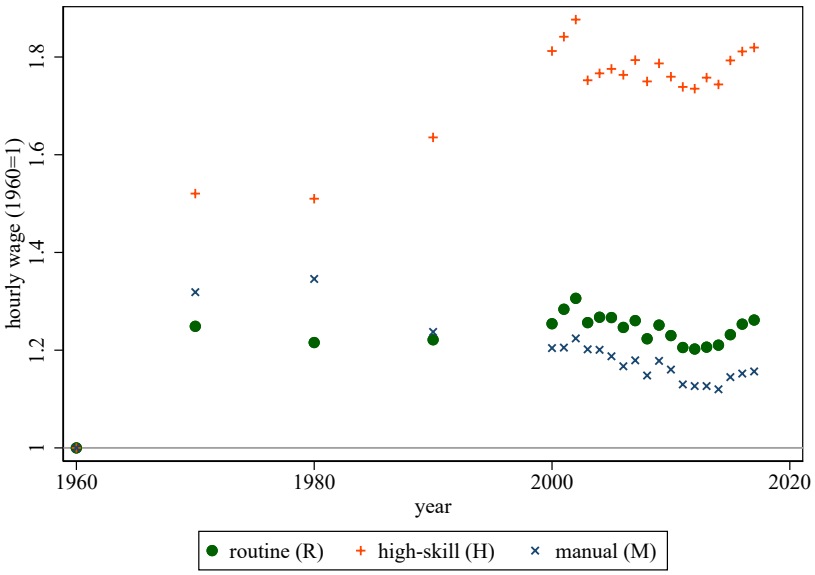
(a) Growth rates of hours worked

(b) Standard deviation of wages



Note: The left panel shows annual growth rates of high-skill and low-skill workers. High-skill workers are defined as workers in non-routine cognitive occupations, low-skill workers as all others. The right panel shows the annual standard deviations of wages for routine and non-routine manual workers as obtained from a truncated maximum likelihood estimation. The occupational classification follows Autor and Dorn (2013). Numbers show a trend from applying an HP filter. Between Census years 1960, 1970, 1990 and 2000, numbers are interpolated.*Source:* ACS, Census and Autor and Dorn (2013).

Figure A3: Hourly wages, ACS



Note: The graph shows the development in hourly wages by occupation group between 1960 and 2017 based on data from the Census and ACS. Hourly wages are normalized to 1 in 1960

Table A1: ICT assets: equipment

asset code	asset name	ICT asset
EP1A	Mainframes	✓
EP1B	PCs	✓
EP1C	DASDs	✓
EP1D	Printers	✓
EP1E	Terminals	✓
EP1F	Tape drives	✓
EP1G	Storage devices	✓
EP1H	System integrators	✓
EP20	Communications	
EP34	Nonelectro medical instruments	
EP35	Electro medical instruments	
EP36	Nonmedical instruments	
EP31	Photocopy and related equipment	
EP12	Office and accounting equipment	
EI11	Nuclear fuel	
EI12	Other fabricated metals	
EI21	Steam engines	
EI22	Internal combustion engines	
EI30	Metalworking machinery	
EI40	Special industrial machinery	
EI50	General industrial equipment	
EI60	Electric transmission and distribution	
ET11	Light trucks (including utility vehicles)	
ET12	Other trucks, buses and truck trailers	
ET20	Autos	
ET30	Aircraft	
ET40	Ships and boats	
ET50	Railroad equipment	
EO11	Household furniture	
EO12	Other furniture	
EO30	Other agricultural machinery	
EO21	Farm tractors	
EO40	Other construction machinery	
EO22	Construction tractors	
EO50	Mining and oilfield machinery	
EO60	Service industry machinery	
EO71	Household appliances	
EO72	Other electrical	
EO80	Other	

Note: The data are retrieved from BEA's detailed fixed assets accounts.

Table A2: ICT and non-ICT assets: intellectual property products

asset code	asset name	ICT asset
ENS1	Prepackaged software	✓
ENS2	Custom software	✓
ENS3	Own account software	✓
RD11	Pharmaceutical and medicine manufacturing	
RD12	Chemical manufacturing, ex. pharma and med	
RD23	Semiconductor and other component manufacturing	✓
RD21	Computers and peripheral equipment manufacturing	✓
RD22	Communications equipment manufacturing	
RD24	Navigational and other instruments manufacturing	
RD25	Other computer and electronic manufacturing, n.e.c.[1]	
RD31	Motor vehicles and parts manufacturing	
RD32	Aerospace products and parts manufacturing	
RDOM	Other manufacturing	
RD70	Scientific research and development services	
RD40	Software publishers[1]	
RD50	Financial and real estate services	
RD60	Computer systems design and related services[1]	
RD80	All other nonmanufacturing, n.e.c.	
RD91	Private universities and colleges	
RD92	Other nonprofit institutions	
AE10	Theatrical movies	
AE20	Long-lived television programs	
AE30	Books	
AE40	Music	
AE50	Other entertainment originals	

Note: The data are retrieved from BEA's detailed fixed assets accounts. As this paper excludes structures, the list contains only equipment capital and intellectual property products. [1]: Excluded due to short time series dimension.

Table A3: Income and wealth by education group, SCF 1998-2013

	Labor income	Capital income	Total income	Net wealth
college graduates	56,443	50	59,300	142,044
non-college graduates	27,873	0	29,707	27,634

Note: All numbers are in 2019 US Dollars and refer to the median household in each education group. Total income is the sum of labor income, capital income and transfer income as defined by Kuhn and Ríos-Rull (2016). *Source:* SCF.

A.2 Details on Data Sources and Dataset Construction

NIPA

The goods and assets price series as well as aggregate consumption measures are constructed using data from the BEA's National Income and Product Accounts (NIPA). For asset prices, we use NIPA Tables 5.5.3, 5.5.4 and 5.5.5, which contain quantities, prices and expenditures for equipment, and Tables 5.6.3, 5.6.4 and 5.6.5 for intellectual property products. We merge the two categories and create aggregate Törnqvist price indices for ICT and non-ICT capital, following the classification used for the Detailed Data for Fixed Assets and described in the main text. For the price series, we use NIPA Tables on personal consumption expenditures by type of product. We use tables 2.4.3U, 2.4.4U, 2.5.4U to construct series for consumed quantities, prices and personal consumption expenditures. To match these with the ICT share measures at the level of IO final goods and services, we use a crosswalk from NIPA PCE lines to UCC codes provided by the BLS. Then we apply our matching of UCC codes to IO final commodities. We construct price indices for the composite ICT- and non-ICT good as described in the main text.

Input-Output Accounts

We use the BEA's Input Output Accounts and focus on the detailed Input-Output tables after redefinitions. We use producer-value tables, which means that the distribution margin (i.e. the cost of wholesaling, retailing and transportation) is modeled as a flow from the distribution industries to the final consumer rather than as a flow from the producing industry to the final consumer. This is the standard approach and is appropriate here as we have no reason to believe that the digitalization content of distribution differs between the different goods and services. We match our 61 BEA industries to the approximately 400 IO product categories.

The Input-Output accounts are presented in a set of different tables, among them *use, make* and *direct requirements* tables. We start from the commodity-by-industry direct requirements table, which shows the amount of each commodity (i.e., intermediate or final good or service) that is required by an industry to produce one dollar of the industry's output. The problem with this table is that the same commodity can be produced by different industries, e.g. ice-cream can be produced by the dairy product manufacturing industry and the ice-cream manufacturing industry. We follow Horrowitz and Planting (2009) in creating a commodity-by-commodity direct requirements matrix that takes the market shares of each industry in producing certain commodities into account. With this matrix, we can calculate the digitalization share of final goods and services as a weighted mean of the digitalization shares of its intermediate inputs and of value added. Note

that all sectors related to housing are matched to the real estate sector.

We pursue an iterative approach. We initially assign to each commodity the digitalization share of the industry that is the ultimate producer of the commodity. Then, we consider all commodities that use this specific commodity as intermediate input. We update the output digitalization share by calculating a weighted mean of the digitalization shares of all inputs and the digitalization share of value added of the final producer. We again assign this share to the inputs used to produce other commodities, and update the output digitalization share again, etc. We continue this procedure until the commodity digitalization shares have converged to fixed values.

For the import share used as a control variable in Table 1, we use IO use tables after redefinitions at producer value. The import share is calculated as the (negative of the) value of imports of goods and services divided by the value of total output.

CEX

The Survey of Consumer Expenditures is the most detailed expenditure survey in the United States, carried out at the household level. Next to providing detailed expenditure data, it also contains a host of demographic and income information on the households. We use the interview survey data of the public-use microdata, which covers the full consumption basket starting from 1990. Sample households are interviewed at a quarterly frequency for four to five consecutive months. We treat each interview round as independent observation. Households report monthly expenditures at the universal classification code (UCC) level for the three months prior to the interview month. Additionally, the data contain quarterly expenses at the level of summary product categories. We consider only households where the head is between 16 and 64 years old and in the labor force. This is because we want to focus on households where labor income is the main source of income. In defining the head, we divert from the CEX convention by making the head the man in mixed couples.

We measure income as total before-tax annual household income. Income is self-reported prior to 2004 (FINCBTAX) and imputed from 2004 onward (FINCBTXM). We only consider households with strictly positive income. We drop the top and bottom 5% of the income distribution in order to mitigate the effect of outliers and top-coding, following Aguiar and Bils (2015). The main income component is labor income. Next to labor income, the income measure includes farm and non-farm business income, social security income, interest on savings accounts or bonds, income from dividends, royalties, estates and trusts and rental income. For renters, we define housing expenditures by rent paid. For home owners, we use self-reported rental equivalence rather than out-of-pocket expenditures for repairs and maintenance and mortgage payments. We also drop financial expenditures such as cash contributions, loans and interest payments, alimony payments and cash gifts, as well as trade-in allowances. We calculate annual expenditures at the level of calendar years. As survey weights refer to the interview year, we need to adjust survey weights for expenditures that relate to the previous year. This concerns households interviewed in January through March. E.g. for a household interviewed in February, the reported expenditures refer to November and December of the previous year, as well as January of the interview year. In the adjustment we follow the instructions published in the Public Use Microdata Users' Documentation.

Before 1990, the CEX only contains spending on summary product categories. For 1984-1989, we follow Jaravel and Lashkari (2024) to impute UCC-level expenditures. They use microdata for 1990 and assume that the expenditure shares for any given income percentile within each of the summary product categories are the same in 1984-1989 as in 1990. Our final sample has between 9,000 to 18,000 observations per year.

The CEX is known for under-reporting of expenditure, in particular by richer households (Aguiar and Bils, 2015; Attanasio and Pistaferri, 2016). This problem has been increasing over time. In

consequence, inequality measures like the total consumption expenditure of rich relative to poor households are biased downwards. For the purpose of this analysis, we consider the spending of households of a specific income bin on different products relative to their total expenditure. As long as rich households underreport all expenditures to an equal extent, our results should not be biased.

We observe only total expenditure on each item, and cannot separate between prices paid and quantities purchased. This is problematic, because households may face different prices for the same good, or choose differently-priced products within the narrow categories of the CEX. This issue has been recognized by the literature on consumption inequality (see Attanasio and Pistaferri (2016) for a discussion) and means that we are potentially overestimating consumption inequality.

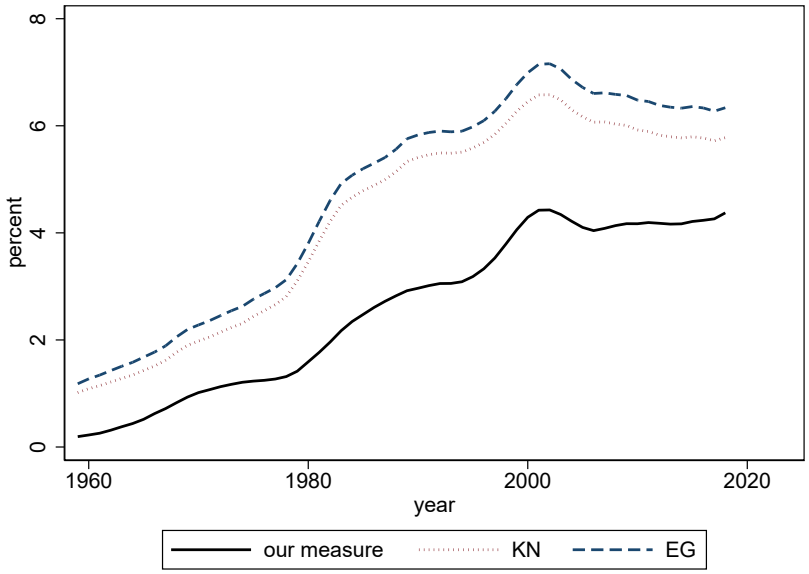
Census and ACS

Most of our empirical measures of employment and wages are constructed using data from the Decennial Census (for the years 1960, 1970, 1980 and 1990) and the American Community Survey (ACS, available annually from 2000). For the years in between, we interpolate the data. In the sample composition as well as the definition of routine and non-routine workers, we follow Autor and Dorn (2013). We keep only individuals between ages 16-64 that are employed and non-institutionalized. Our key measures are the number of hours worked per year and the hourly wage. The number of hours worked is the number of hours usually worked during a week times the number of weeks worked within a year. As income is top-coded, we replace income in the top bracket with 1.5 times the actual value and then divide by the number of hours worked to obtain the hourly wage. We report all Dollar values in constant 1999 US Dollars. We create a consistent occupational classification using the crosswalks provided on David Dorn's website. We classify the following occupations as routine: sales and administrative support, precision production and craft workers and machine operators. We classify as non-routine cognitive workers managers, professionals and technicians. Non-routine manual occupations are those working in construction and mining, services, mechanics, repairers and transportation. We drop individuals working in the agricultural sector or the military. Out of the total of 318 occupations, 210 are routine or non-routine manual.

A.3 Comparison of ICT share to related measures

Both Karabarbounis and Neiman (2019) and Eden and Gaggl (2018) calculate the ICT share in value-added, but using a slightly different definition of ICT assets. Their set of ICT assets is broader than ours, also including medical and nonmedical instruments, photocopy equipment, office and accounting equipment and communications equipment. In contrast, Karabarbounis and Neiman (2019) do not include intellectual property products that relate to computer and electronic manufacturing and software publishers. Computers and software, which are at the core of our measure, have seen stronger price declines than other assets such as medical and nonmedical instruments, photocopy equipment or office and accounting equipment. None of the other two papers takes quality changes in ICT capital into account. For the sake of comparability, we have computed their measures analogously to ours, i.e. adjusting for quality improvements. Figure A4 shows ICT shares for the aggregate US economy. Due to the broader ICT definition in Karabarbounis and Neiman (2019) and Eden and Gaggl (2018), their ICT shares are larger. However, they exhibit similar time trends.

Figure A4: ICT share comparison



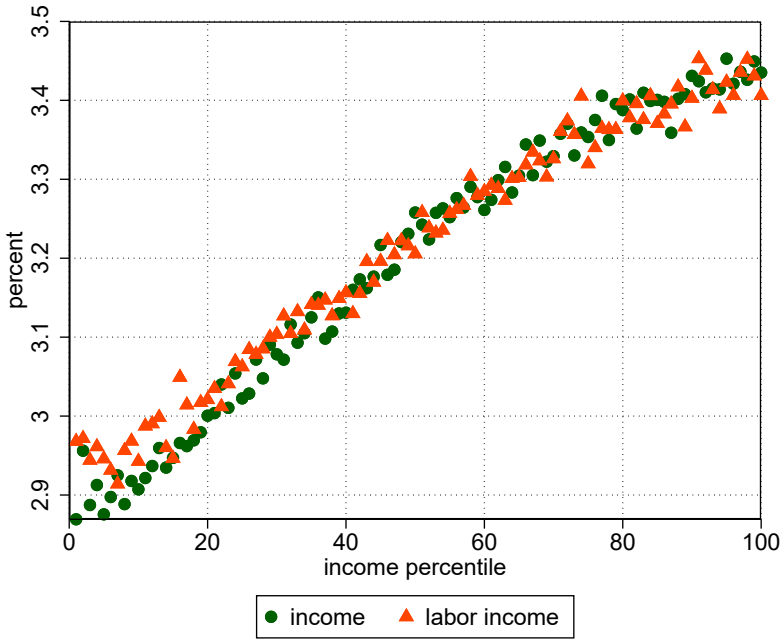
Note: The figure compares the ICT share in value-added of this paper with Karabarbounis and Neiman (2019) and Eden and Gaggl (2018). For the sake of comparability, all three measures are calculated using quality-adjusted asset prices. Note that the numbers for our ICT share differ slightly from the numbers shown in Figure 2 (a) since we calculate a weighted average here, whereas Figure 2 is an unweighted average across industries. *Source:* BEA and own calculations.

A.4 ICT share in the Consumption Basket: Extensions and Robustness

Sorting Households by Labor Income

As an alternative to grouping households by total income, we group households into percentiles of labor income, dropping households with non-positive labor income. As Figure A5 shows, it makes almost no difference whether we consider percentiles of total income or labor income, the only difference being the lowest percentiles.

Figure A5: ICT share along the labor income distribution

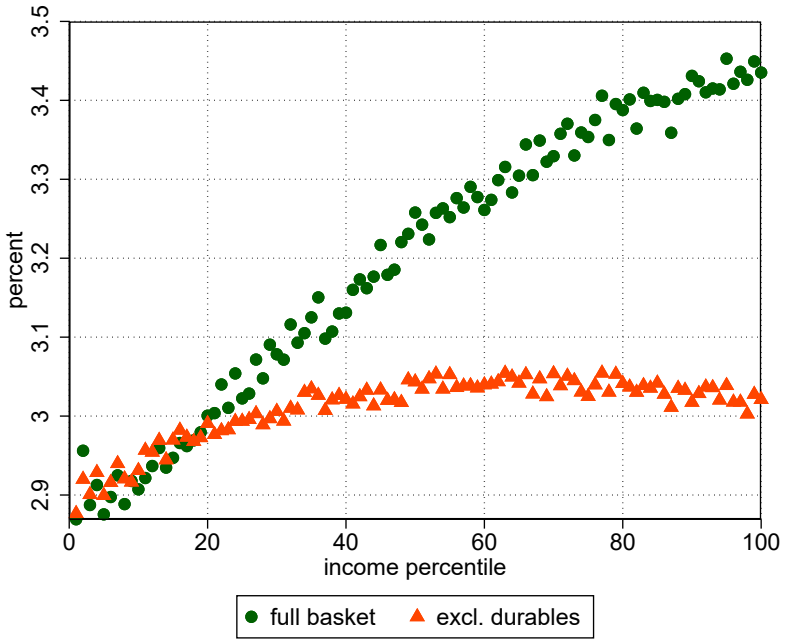


Note: The graph shows the ICT share of the consumption basket comparing percentile of total household income with percentiles of labor income. *Source:* BEA, CEX and own calculations.

Non-durable Goods

Consumption in most macroeconomic models can best be understood as consumption of non-durable goods, therefore we exclude in this extension spending on durable goods such as housing, vehicles or education. Figure A6 is the equivalent to Figure 5 but excluding durable goods. This excludes around 35% of expenditures. Among those dropped are expenditures for housing, household equipment, vehicle expenses, health and education. We follow Aguiar and Bils (2015) in the non-durable/durable classification. The ICT share still increases along the income distribution, but the pattern is significantly weakened. Durable goods are an important driver of differences in ICT share across households.

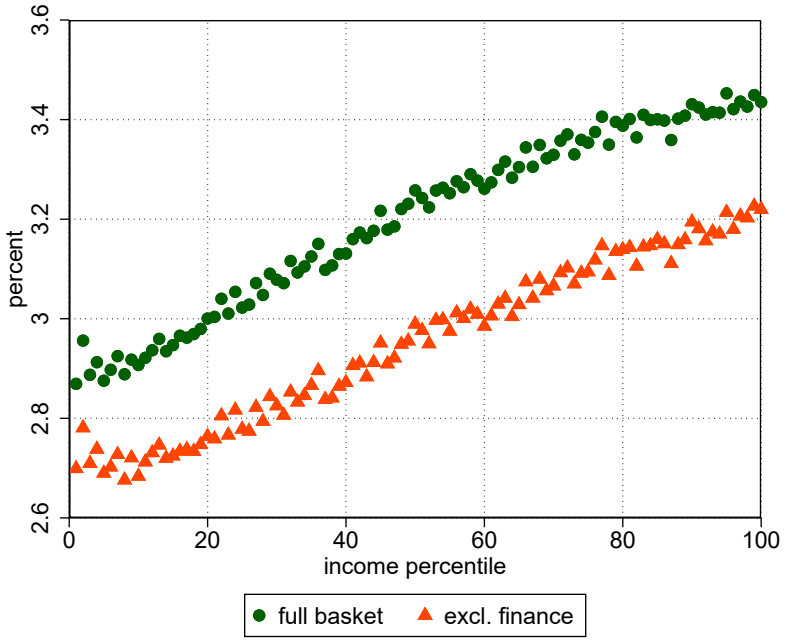
Figure A6: ICT share along the income distribution, only non-durable goods



Note: The graph shows the ICT share of the consumption basket by percentile based only on non-durable goods.*Source:* BEA, CEX and own calculations.

Excluding Finance and Insurance

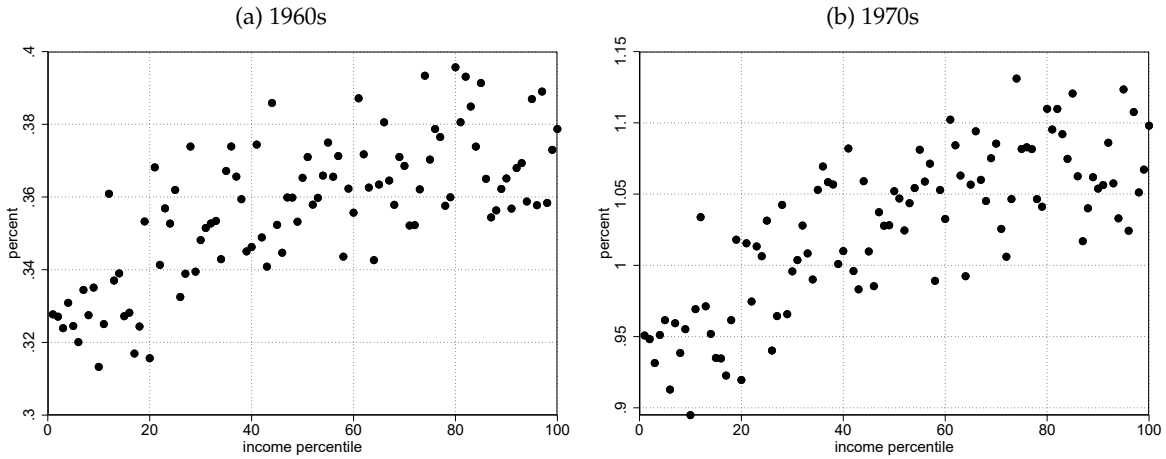
Figure A7: ICT share along the income distribution, excluding finance and insurance



Note: The graph shows the ICT share of the consumption basket by percentile for the full sample period including and excluding finance and insurance. *Source:* BEA, CEX and own calculations.

Historical Consumption Data

Figure A8: ICT share along the income distribution, historical data



Note: The graph shows the ICT share of the consumption basket by income percentile for the 1960s and 1970s.*Source:* BEA, CEX, Jaravel and Lashkari (2024) and own calculations.

We use historical data constructed by Jaravel and Lashkari (2024) for the 1960s and 1970s. These data are based on several imputations: first, between the years with available data (1960, 1972, 1984). Second, for detailed consumption categories (UCCs), which were not reported in the early waves. Jaravel and Lashkari (2024) use expenditure on summary categories and assume that expenditure shares within these categories were the same as in 1990. As a result, the data are much noisier. The early data do not contain some of the expenditure categories with the highest ICT shares such as computers and TVs, and in consequence the overall level of the ICT share is much lower. However, the positive income gradient is also observable in the 1960s and 1970s.

Demographic Factors

In this subsection, we study what role the demographic factors age and education play for our results. Household income tends to increase with age, which is why we might be concerned that rather than capturing an increasing ICT share of consumption along the income distribution, we capture higher spending shares on digital goods among the age distribution. Similarly, income also tends to be higher for households with college education, who at the same time might consume more digitally. Whether a household is located in an urban or rural area might also place a role for digital consumption patterns, while at the same time it correlates with income. To refute these concerns, we first run regressions of the ICT share on income, controlling for the age and the education of the household head and for whether the households lives in a urban or rural area. Then, we show Figure 5 separately for young households (below age 32), middle-aged households (aged 32-47) and old households (above age 47), households with and without college education, and urban and rural households.

Table A4 shows regressions of ICT share on income and demographic controls at the level of income percentiles. The demographic controls are measured as averages across all households within the income percentile, where age and education are for the household head. Thus, age is the average age of the household head, educ the average level of education (where at the individual level, 1 means no high-school degree, 2 means high-school but no college degree, and 3 college degree), and urban the share of households living in urban areas. The coefficient on income is highly significant and remains significant once demographic control variables are added.

Table A4: Regression of ICT share on income and age

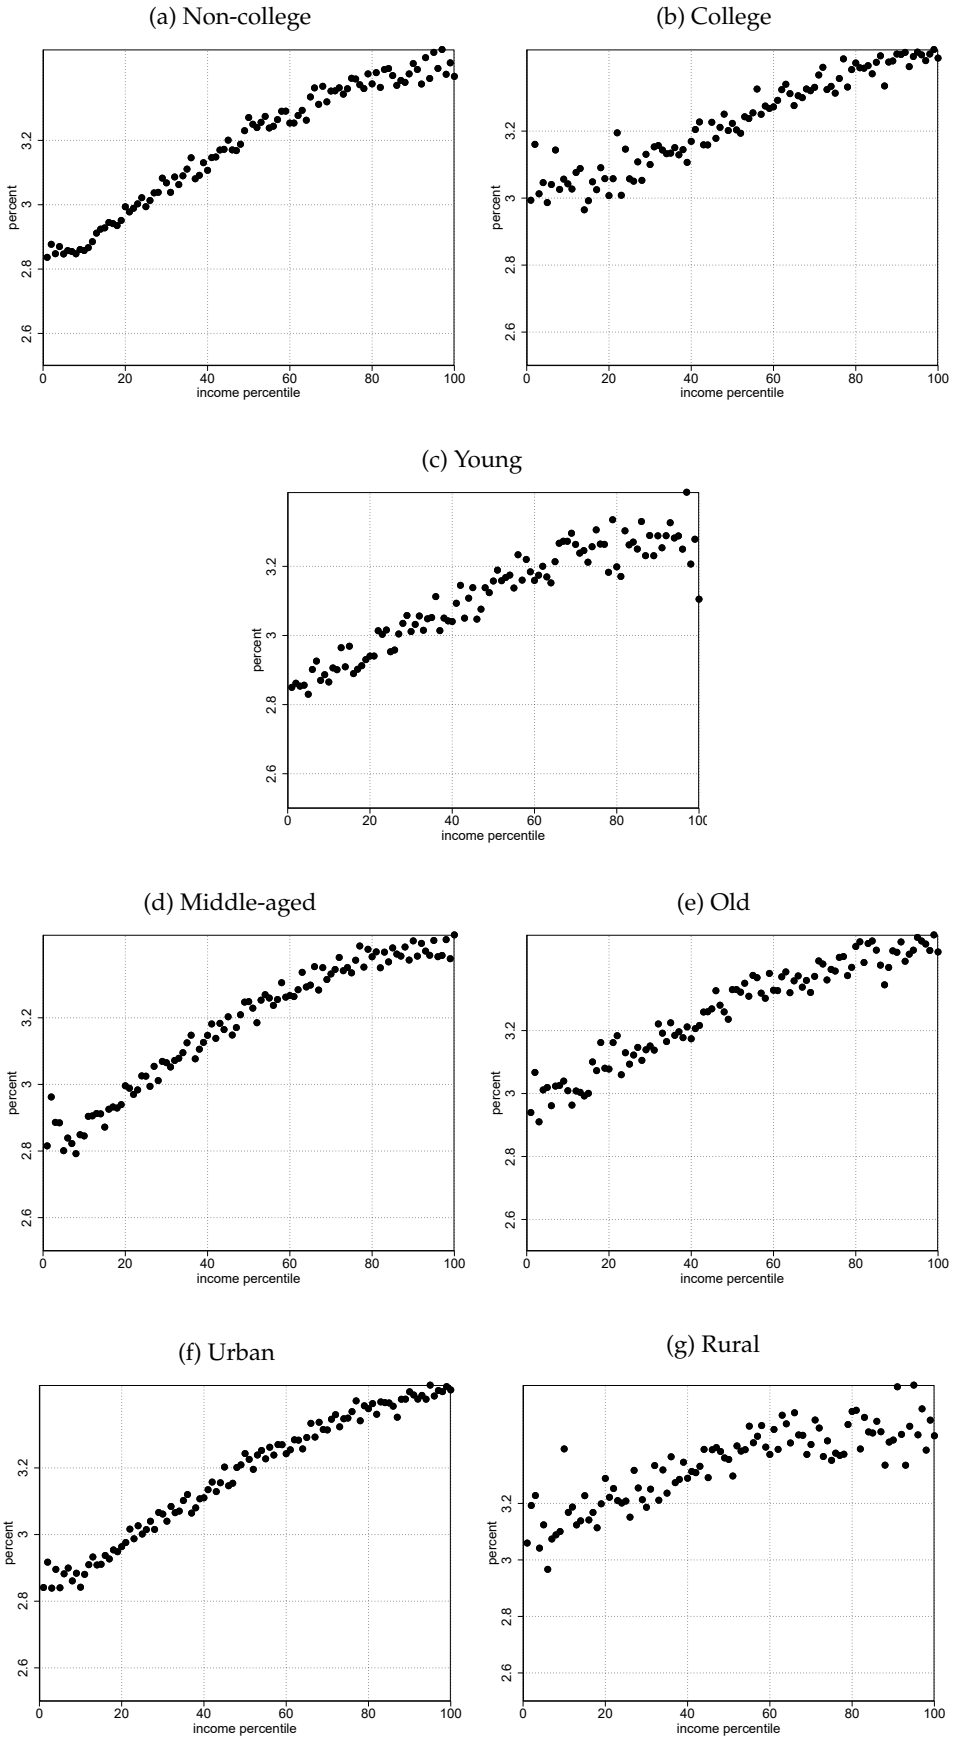
	(1)	(2)	(3)	(4)	(5)
income	0.483***	0.323***	0.223***	0.492***	0.174***
	(0.007)	(0.010)	(0.013)	(0.007)	(0.013)
age		2.617***			1.747***
		(0.134)			(0.137)
educ			48.092***		42.639***
			(2.089)		(2.235)
urban				-17.928**	-61.737***
				(7.094)	(6.588)
Constant	181.311***	84.531***	87.143***	197.328***	88.354***
	(1.329)	(5.108)	(4.272)	(6.476)	(7.209)
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	3,400	3,400	3,400	3,400	3,400
R-squared	0.958	0.962	0.963	0.958	0.966

Note: Income is in thousands of constant 2015 USD. ICT share is in %. Education and age refer to the household head. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

As expected, the magnitude becomes smaller when adding age and education. In contrast, controlling for the urban share makes the income gradient slightly steeper. Including all demographic controls in column (5), the coefficient on income is still positive and significant, but only about one third of the size of column (1).

In Figure A9, the ICT share increases with income across all demographic subgroups. The non-college group shows a larger income gradient than the college group, so the difference between the bottom and top percentile is more pronounced for non-college households, suggesting that they are a more heterogeneous group than college-educated households. The median for both groups is very similar (2.99 vs. 2.94). The ICT share is slightly higher for older households than for middle-aged and younger ones. The median is 2.88 for the young compared to 3.04 for the old. The rural sample is much smaller than the urban sample and therefore exhibits more dispersion. The median in the rural sample is – perhaps counterintuitively – higher than in the urban sample (2.96 vs. 3.96).

Figure A9: ICT share along the income distribution - subsamples



Note: The graph replicates Figure 5 for households (a) without college degree, (b) with college degree, (c) aged below 32, (d) aged between 32-47, (e) aged above 47, (f) living in urban areas, (g) living in rural areas. For the ease of exposition, we only show the average across the whole time period 1984-2017. *Source:* BEA, CEX and own calculations.

A.5 Relative Price Changes: Extensions and Robustness

Additional Price Regressions

Table A5 shows several robustness checks for Table 1: In columns (1) and (2), we include fixed effects at the level of IO goods and services instead of major industries. In columns (3), (4) and (5) we use as explanatory variable the beginning-of-decade ICT share instead of end-of-decade ICT share. In columns (6), (7) and (8), we include a quadratic ICT share. The baseline results of Table 1 are robust to each of the specifications. Notably, the ICT share seems to have a non-linear effect on price changes, and is less negative at higher ICT shares. The net effect would turn positive from an ICT share of 42.0, which is however outside of the bounds of our sample. The highest ICT share in our sample is 33.0, which is found for the software publishing industry (511200).

Table A5: Price changes and ICT share

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
VARIABLES	60-17	60-17	67-17	67-17	67-17	60-17	60-17	60-17
ICT share	-0.0743***	-0.0804**				-0.151	-0.184**	-0.294***
	(0.00953)	(0.0331)				(0.140)	(0.0726)	(0.111)
L.10 ICT share			-0.205***	-0.163***	-0.182***			
			(0.0643)	(0.0392)	(0.0458)			
ICT share^2						-0.000562	0.00170	0.00495
						(0.00424)	(0.00220)	(0.00320)
Constant	0.0394	0.0226	0.328	0.212	0.377	0.298	0.270	0.207**
	(0.0401)	(0.185)	(0.327)	(0.205)	(0.237)	(0.338)	(0.212)	(0.0991)
Observations	3,912	3,912	3,275	3,275	3,275	3,912	3,912	3,912
R-squared	0.531	0.533	0.091	0.334	0.338	0.102	0.344	0.351
industry FE	no	no	no	yes	yes	no	yes	yes
IO FE	yes	yes	no	no	no	no	no	no
decade FE	no	yes	no	no	yes	no	no	yes

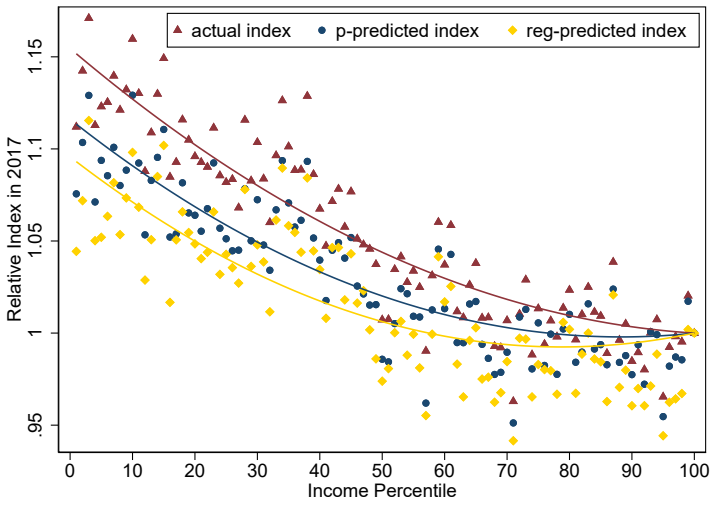
Note: The table shows the regression of annualized decadal price changes on the ICT share at the end or beginning of each decade. The unit of observation are product categories. All variables are measured in %. Standard errors clustered at the level of final goods and services in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Alternative Price Indices

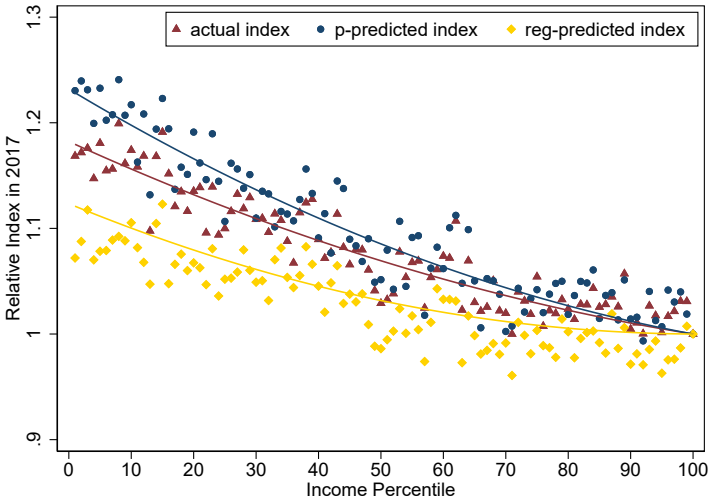
Figure A10 replicates the result of Figure 7 but using two different price indices: the Laspeyres index and the Paasche index. It is a well-known fact that the Laspeyres index, due to its reliance on period t-expenditure shares, tends to overstate inflation, the Paasche index, relying on period t+1-expenditure shares, understates inflation. Jaravel (2019) showed that inflation inequality across the income distribution looks similar for all three indices. Figure A10 confirms this pattern for the longer 1984-2017 period while also showing that the effect of digitalization on prices is similar for Törnqvist, Laspeyres and Paasche.

Figure A10: Inflation and ICT, alternative price indices

(a) Laspeyres price index



(b) Paasche price index



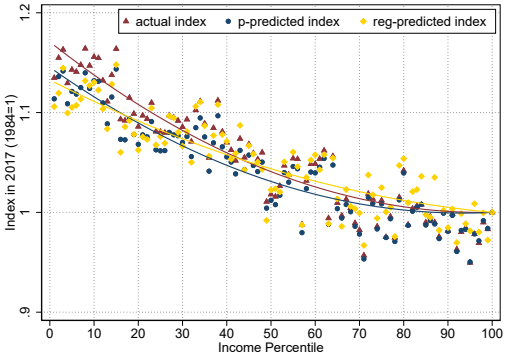
Note: The top graph plots Laspeyres price indices, the bottom graph Paasche price indices. Both price indices are shown for 2017, indexed to 1 in 1984 and relative to the price in the highest income percentile. The red triangles indicate the actual data, whereas the blue circles and the yellow diamonds indicate ICT-adjusted indices following the same computation as in Figure 6. The blue ICT-adjusted Paasche index indicates that ICT-intensive products would benefit the poor, but this is due to stronger substitution among the poor, which gets a disproportionate weight. *Source:* BEA, CEX and own calculations.

Demographic Factors

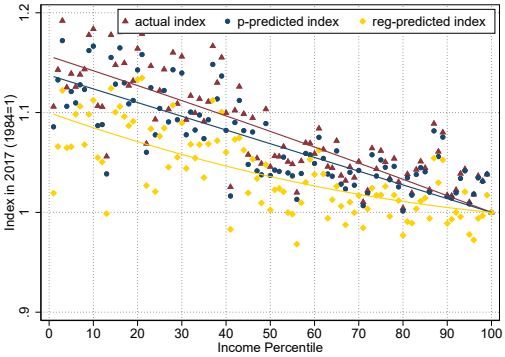
Figure A11 replicate Figure 7 for the five demographic subgroups discussed in Appendix Section A.4: Young households (below age 32), middle-aged households (aged 32-47) and old households (above age 47), households with and without college education, and households living in urban vs. rural areas. Due to the low number of observations in some subgroups (in particular the rural and the young subgroups) there is a lot of dispersion in the data, which makes it difficult to interpret.

Figure A11: ICT share along the income distribution - subsamples

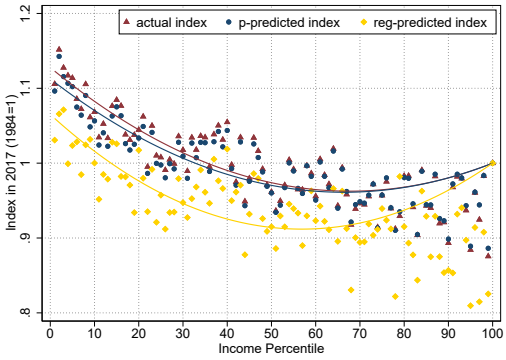
(a) Non-college



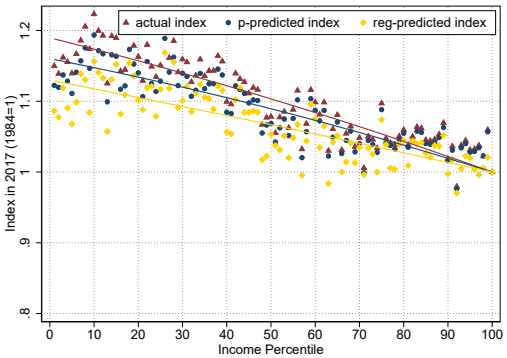
(b) College



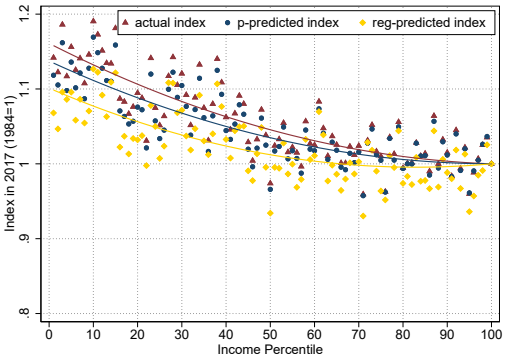
(c) Young



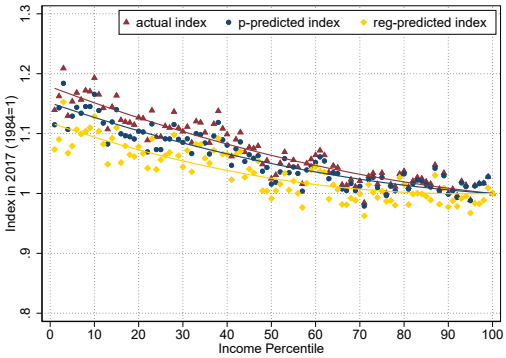
(d) Middle-aged



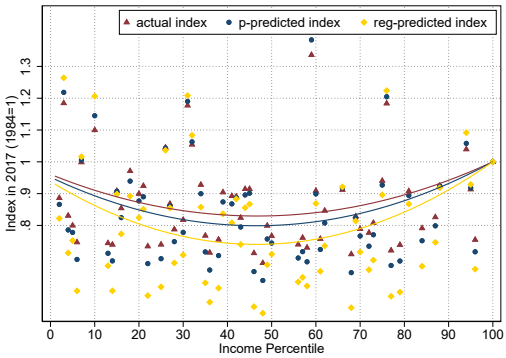
(e) Old



(f) Urban



(g) Rural



Note: The graph replicates Figure 7 for households (a) without college degree, (b) with college degree, (c) aged below 32, (d) aged between 32-47, (e) aged above 47, (f) living in urban areas, (g) living in rural areas.
Source: BEA, CEX and own calculations.

A.6 Alternative Digitalization Measures

Input- vs. Output-based Digitalization Measures

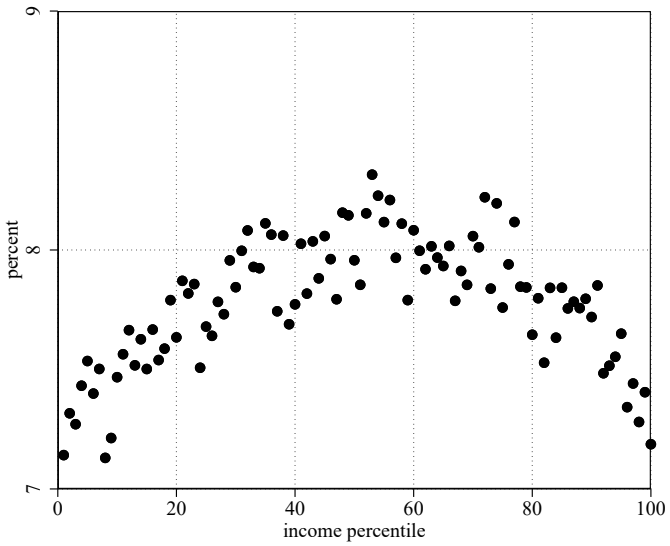
In the main text, we define digitalization based on the inputs used in the production process. A good is classified as digital according to the extent to which it has been produced with digital assets like computer hardware or software. An alternative would be an output-based measure, sorting final consumption goods into digital or non-digital goods. Different attempts have been made to classify final consumption goods. We rely on the BEA's Digital Economy Satellite Account, which represents the BEA's approach to calculating the contribution of digital goods- and service-producing industries to U.S. GDP ("digital economy"). The BEA defines the digital economy as comprising three types of goods and services: digital infrastructure, which primarily refers to ICT, e-commerce and priced digital services, i.e. services related to computing and communication that require a fee (BEA, 2021). The BEA provides data for the digital economy's value added at the industry level. We use the share of digital value-added in total value-added as output-based ICT share measure. Since the industry classification is the same as the one we use to construct our output-based measure, we can easily compare it to our input-based measure.

The BEA Digital Economy Satellite Account does not (yet) attempt to estimate the value of free digital goods and services. It is sometimes approximated by the advertising revenue generated by the companies offering it, however, this will insufficiently reflect the consumer surplus generated by these products (Brynjolfsson et al., 2020). Brynjolfsson et al. (2019) and Coyle and Nguyen (2020) present an alternative approach to estimate the value of free digital goods. They carry out choice experiments to estimate consumers' reservation prices, asking them about the monetary compensation at which they are willing to give up access to a digital good such as Google Maps, Facebook or Wikipedia. These studies show that the consumer surplus generated from free goods can be substantial. However, the research in this area is still in its infancy and the existing literature does not study differences between high- and low-income households systematically. Therefore, it is not possible to assess the bias created by ignoring the consumer surplus of free digital goods.

The correlation between our input-based and the BEA's output-based measures is 0.78. While both measures rely highly on the ICT sector (Naics 334* (excl. 3345), 5112, 5191, 5415 and 518), the following categories have a high ICT share according the input-based measure, but not the output-based measure: navigational, measuring, electromedical and control instruments (3345), financial services and insurance carriers (52*), newspaper and book publishers (5111). These tend to be consumed proportionately more by high-income households. Categories with a low input ICT-intensity but high output ICT-intensity are in particular radio and television broadcasting (5151) and telecommunications (517). The expenditure share of these categories is higher for low-income households. Due to these differences, the ICT share along the income distribution looks different when using the output-based definition. Figure A12 shows the digital share of the consumption basket along the income distribution for the output-based measure. According to this measure, the ICT share is hump-shaped along the income distribution, with middle-income households having the highest ICT share, whereas top and bottom are very similar to each other.

This underlines that it matters how digital goods are defined. Choosing one measure or the other might lead to different implications about who benefits and who loses from digitalization. We have opted for an input-based measure in this paper since we are interested in comparing income and price effects of changes in production processes, but we acknowledge that there could also be benefits in studying output-based measures, and that the findings would be complementary to ours.

Figure A12: ICT share along the income distribution, output-based classification



Note: The graph shows the ICT share of the consumption basket by percentile for 2005-2017 period for the output-based series constructed based on the BEA's Digital Economy Satellite Account. *Source:* BEA, CEX and own calculations.

Digitalization and Time Use

We supplement our main analysis with data from the American Time Use Survey (ATUS). ATUS documents how Americans allocate their time across different activities. It has been running since 2003 and is based on a subsample drawn from the Current Population Survey (CPS) of between 10,000 to 20,000 respondents. Within a 24-hour period, the respondents document their primary activities within 15 minute windows, choosing among more than 400 different activity categories. We exclude from our analysis the time when individuals are working and sleeping as well as the use of government services.

To develop a concordance that matches ATUS activity codes to IO categories, we follow Fang et al. (2021). However, while Fang et al. (2021) match broad ATUS categories to broad CEX expenditure categories, our concordance is more disaggregate.

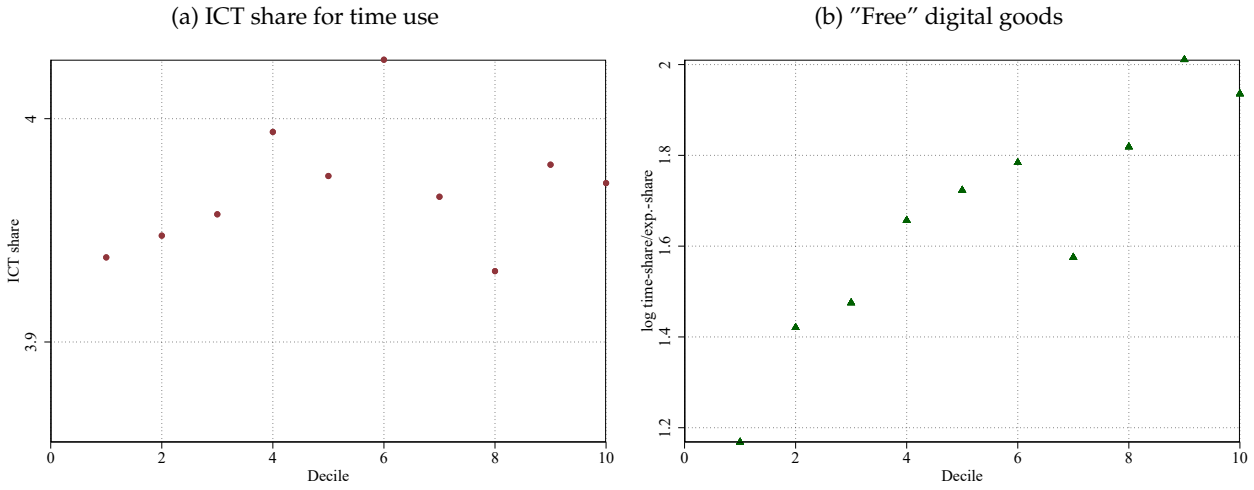
When matching all activities to their corresponding expenditures, there is a near zero correlation between expenditure and time-shares for all income groups. This is illustrated by Table A6 which focuses on the richest and poorest 10 percent. This implies there is no overall pattern suggesting time use and expenditures are substitutes or complements. This is also in line with the fact that the average ICT share is similar for both time use and consumption expenditures, as illustrated by comparing Figure 5 with A13 (a). Among categories that are substitutes (high time use and low expenditures or vice versa) are activities related to housework. Those are slightly less ICT intensive and more important for the poor.

Table A6: Correlation between time share and expenditure share

Time Period	Top 10%	Bot 10%
2003-2010	0.05	0.04
2011-2017	0.01	0.04

Sources: BEA, CEX, ATUS and own calculations

Figure A13: Time use across income deciles



Note: Graph (a) shows the ICT share of the overall time use by income decile for different sub-periods between 2003 and 2017. Graph (b) shows the log ratio of time-expenditure-shares for "free" digital products and services, or more precisely digital activities with zero marginal costs. Free digital products and services are defined as ATUS activities computer use, e-mail services, playing games and listening to music. Average between 2003-2017. *Source:* BEA, ATUS, CEX and own calculations.

Next, we compute the average ICT share of time use of an individual by summing up the time-share weighted activities multiplied with their ICT shares. We track the time share of activities since 2003 and compute the average ICT time use intensity across income deciles[12]. Analogous to the ICT share of expenditures in Figure 5, Figure A13 (a) shows the ICT share of time use along the income distribution. The ICT share in time use, while noisy, moves similar to the expenditure share between 3.9 and 4.1% and features a slight upward trend along the income deciles.

ATUS partly captures free digital products and services in the categories computer use, e-mail services, playing games and listening to music. The notion of "free digital good" relates to near zero marginal costs, once these goods are accessible. While the CEX allows measuring the entry costs for these goods, ATUS shows how much time households dedicate for these activities. We compute the log ratio between time shares and expenditure shares for each income decile, to assess whether a group dedicates relatively more time to free digital goods. The ratio is upward sloping in Figure A13 (b) which implies that rich households spend relatively more time on 'free' activities relative to their expenditures than poor households. This could suggest that entry costs (such as buying a computer or having internet access) are relevant, which is problematic for poor households.

The results based on time use are in line with those based on expenditure data. This implies that even though we disregard some parts of consumption when focusing on expenditure, we still get a fairly accurate picture of the differences in ICT shares of the consumption baskets of rich and poor households. Considering the evidence on free digital goods and home production, the digital inequality based on expenditure data is likely to be understated. In Section 3, we assume that all consumption shows up in form of purchases of goods or services at a market price, which therefore constitutes a conservative estimate of the impact digitalization has on inequality.

[12]The household income data of ATUS are reported in few income bins, which is why we consider deciles here.

B Model Appendix

B.1 Sectoral Classification

In the main part of the paper, we sort IO industries into ICT- and non-ICT sector based on their factor inputs, i.e. whether the industries *use* ICT capital. Below we explore two alternative classifications: (1) classification based on whether industries *produce* ICT capital, (2) a categorization following the structural change literature into agriculture, manufacturing and services (Herrendorf et al., 2014) or alternatively into agriculture/manufacturing, high-skill services and low-skill services (Buera and Kaboski, 2012).

Classification Based on Capital Flows

Table 11 in NIPA's 1997 Capital Flow Table shows the input-output commodity composition of private equipment and software expenditure categories. We focus specifically on the assets which we classify as ICT assets in Table A1. Computers and peripheral equipment are produced by BEA industries 3340 (computer and electronic product manufacturing; about 94% of the total value of production) and 5415 (computer systems design and related services, the remaining 6%). Software is produced by 5110 (Publishing industries, except internet, including software; 33% of the total value of production) and by 5415 (66%), the rest falls on non comparable imports. About 80% of the value of non-ICT assets is produced by non-ICT industries. Industries 3340, 5415 and 5110 are a subset of our five ICT-using industries of the main part.

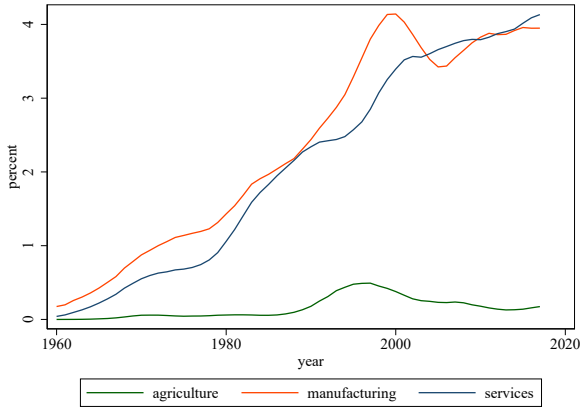
Classification Following the Structural Change Literature

Agriculture – manufacturing – services: Manufacturing and services have an almost identical ICT share (Figure B1 a). Whereas among the industries with the highest ICT share are mostly service industries (especially 5110, 5230, 5415, 5140), the industry on top of the list lies within manufacturing (3340, computer and electronic products manufacturing, with an average ICT share of 17.0%). This distinction is therefore not very informative about digitalization.

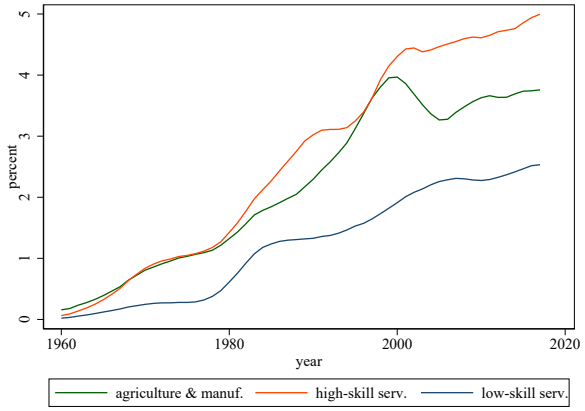
Agriculture and manufacturing – high-skill services – low-skill services: Like the definition in the main part, the high-skill vs. low-skill services distinction aims at the factor inputs, but in contrast to our definition, it focuses exclusively on labor inputs. Buera and Kaboski (2012) classify services as high- and low-skill based on the share of college graduates among the workers employed in the sector. They argue that growth in services is driven by the movement of consumption into more skill-intensive output. This is actually very similar to our framework: high-skill workers are complements to ICT capital, whereas routine workers (a subset of low-skill workers) is substitutable. In fact, all of the high-ICT service industries 5110, 5230, 5415, 5140 are high-skill services. Not surprisingly, the ICT share of high-skill services is higher than that of the other two categories in Figure B1 b). Thus, the rise in ICT-intensive production goes along with the rise in skill-intensive services. A major difference, however, is the role played by manufacturing, or specifically, the role played by the computer and electronic products manufacturing industry. Another difference is that high-skill services according to Buera and Kaboski (2012) include health and education, which are low-ICT.

Figure B1: ICT shares of alternative industry composites

(a) Agriculture – manufacturing – services



(b) High-skill and low-skill services



Note: The graph shows the ICT share of composite sectors. The left graph classifies sectors into agriculture, manufacturing (including also mining and construction) and services. The right graph classifies sectors into agriculture/manufacturing, high-skill services (comprising BEA industries 50*, 51*, 52*, 53*, 54*, 55*, 61*, 62* and 72* as well as 4810) and low-skill services (comprising all other services). *Source:* BEA and own calculations.

B.2 Calibration Algorithm

We calibrate the model in the following way:

1. Global search: We solve the model for 20,000 random draws from a holton set of the nine parameter values $\eta, \epsilon, \phi_1, \phi_2, \gamma_1, \gamma_2, \nu, \theta, \rho, g_{a1}, g_{a2}$ within permissible bounds. We define the loss function as the weighted sum of squared errors in the set of data moments (scaled by the corresponding moment). This means that we minimize the percentage error of moments. We choose a weighting matrix that over-weights by a factor of ten the following, more relevant, moments: the relative price, the skill premium, the wage premium of routine relative to manual workers, and expenditure share differences between the top 1 and bottom 1 percent.

2. Local search: We solve the model for 10.000 random draws in the vicinity of the 10 best solutions, obtain a new set of 10 best solutions and repeat this procedure three times. From the resulting final set of best solutions, we select the one where the loss function is smallest.

B.3 Link between Parameters and Moments

Several moments link directly to one or more parameters. The difference in the expenditure shares of the ICT-intensive good between two income groups is used as a moment to inform ρ. From eq.(20),

$$\frac{P_1 C_{1M}}{e_M} / \frac{P_1 C_{1H}}{e_H} = \left(\frac{e_H}{e_M}\right)^{\rho} \Delta_M$$

where we use as empirical approximation the expenditure shares of the top and bottom 10% of the income distribution as derived in Section 2.

θ is pinned down by the increase in the expenditure share of ICT-intensive goods over time. At

any t, s,

$$\frac{P_{1t}C_{1t}}{e_t} / \frac{P_{1s}C_{1s}}{e_s} = \frac{\left(\frac{P_{2t}}{e_t}\right)^{\rho}\left(\frac{P_{1t}}{P_{2t}}\right)^{\theta}\Delta_t}{\left(\frac{P_{2s}}{e_s}\right)^{\rho}\left(\frac{P_{1s}}{P_{2s}}\right)^{\theta}\Delta_s}$$

The relative quantities consumed of good 1 and good 2 are informative for all three parameters of the utility function:

$$\frac{C_1}{C_2} = \frac{\nu \frac{e_o}{P_1}\left(\frac{P_2}{e_o}\right)^{\rho}\left(\frac{P_1}{P_2}\right)^{\theta}}{\frac{e_o}{P_2}\left(1-\nu\left(\frac{P_2}{e_o}\right)^{\rho}\left(\frac{P_1}{P_2}\right)^{\theta}\right)}$$

where $o = H, R, M$.

To pin down the parameters of the production function, we choose moments that are informative about the relative use of the inputs R, M, H, ICT and K (which can move freely across sectors).

The skill premium of H over R is given by

$$\frac{W_H}{W_R} = \frac{\gamma_i}{1-\gamma_i}\frac{R_i^{\frac{1}{\eta}}}{H^{\frac{1}{\epsilon}}} AW^{\frac{\epsilon-1}{\epsilon}-\frac{\eta-1}{\eta}}$$

Using log difference, γ is not time varying and the premium is therefore strongly influenced by the two elasticities.

The wage premium of R over M additionally incorporates information about the two elasticities:

$$\frac{W_{Rt}}{W_{Mt}} = \frac{AW_{it}^{\frac{\epsilon-1}{\epsilon}}(1-\gamma_i)(1-\psi_i)\phi_i M_{it}}{\psi_i R_{it}^{\frac{1}{\eta}} AW_{it}^{\frac{\eta-1}{\eta}} SW_{it}^{\frac{\epsilon-1}{\epsilon}}}$$

Considering log differences, the two parameters that remain to match this moment are ϵ and η.The time series for W_{Rt}/W_{Mt} is used to match the employment share of routine and manual workers over time.

To match input factors and their returns, we choose the following additional moments: the ICT income share in the aggregate economy, the return on ICT capital and the labor share, which is informative about the use of capital vs. labor in the production process. Koh et al. (2020) show that the decline in the labor share is fully explained by the rise in intellectual property products (whose rents are attributed to capital income by the BEA). A major share of intellectual property products is software. Karabarbounis and Neiman (2014), Martinez (2021), Eden and Gaggl (2018) and Acemoglu and Restrepo (2018) also argue that automation is an important driver of the labor share. Hémous and Olsen (2022) choose a similar calibration strategy to ours where they also require that automation explains the declining labor share and the increasing skill premium fully. Finally we match the evolution of the relative price from Figure 8. This is decisive for the overall strength of the price effect and also helps to inform other parameters in the production function. GDP per capita growth rates and sectoral value added shares then pin down the contribution of TFP growth.

B.4 Validation

Sector-specific Non-targeted Moments

Table B1 shows how the model matches additional non-targeted moments mostly relating to the sectoral composition. Compared to the ICT increase in the overall economy, the increase in the ICT intensive sector is smaller, both in the model and in the data. Furthermore, the increase in the share of hours worked in the ICT intensive sector of high-skill and low-skill workers is very well matched as well. The model slightly over predicts the increase in hours shares for both skill groups, which is probably due to the fact that workers can switch sector at no cost and at any time in the model, whereas in reality there could be substantial frictions connected with switching industry, such as acquiring new skills in the form of additional education or on-the-job training, building a new professional network, etc.

Table B1: Validation: Model and data increases by factors in sector 2

Moment	Data	Model
ICT intensity	3.91	3.46
Share hours H	1.06	0.89
Share hours L	0.21	0.22

Note: All moments are percentage changes between 1960 and 2017. Sources for the data: ACS and BEA.

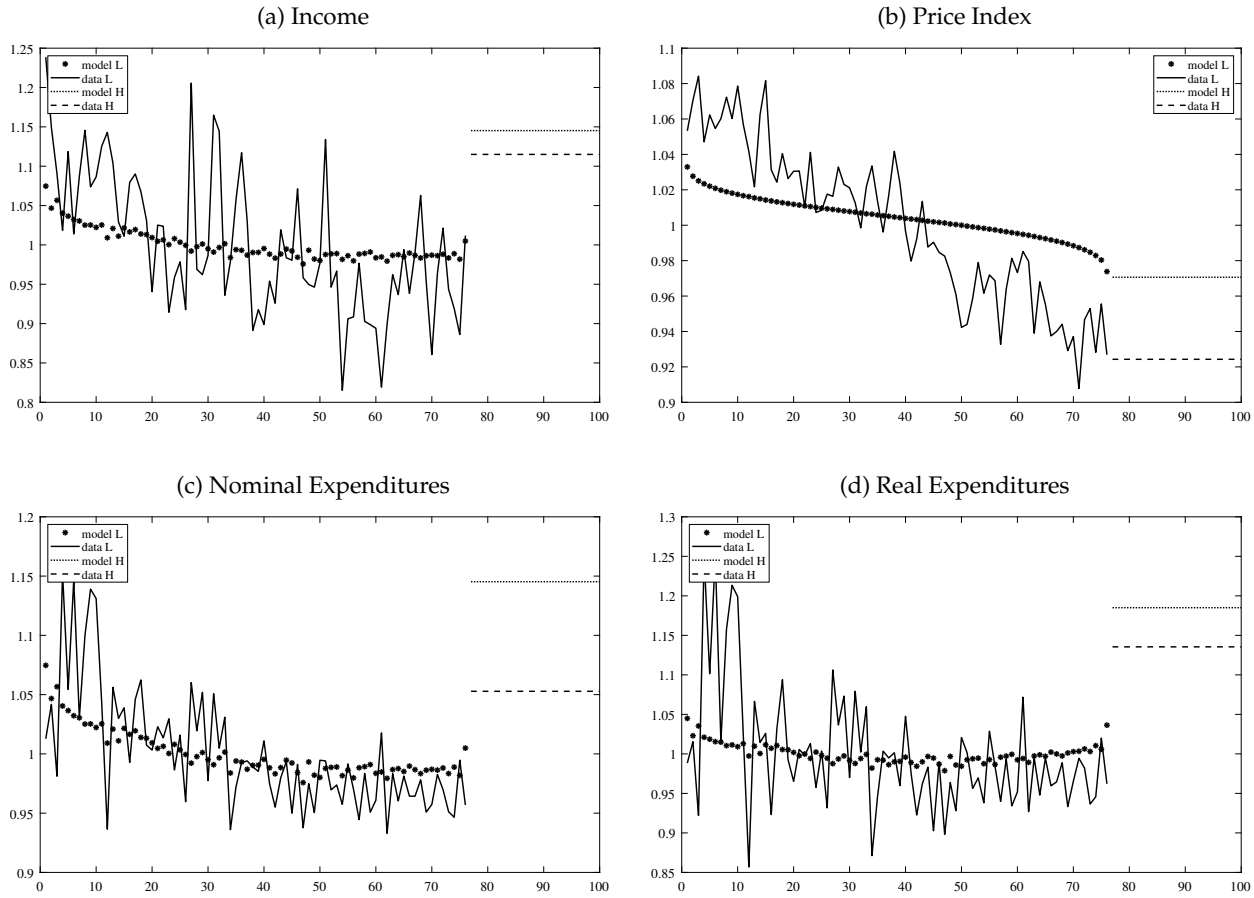
Data-Model Comparison along the Income Distribution

In Figure B2, we plot the model predictions against the data for income, prices and consumption expenditures. To highlight differences across the income distribution, we normalize the mean across low-skill percentiles to 1 for both data and model, thus interpreting each value relative to the average change. Since high-skill workers are homogeneous in the model, the values will appear flat on the right-hand side of the graphs, and there is a discrete jump between percentile 77 and 78. In reality, the change should be more gradual, and the values for percentiles 78-100 should be interpreted as averages across the top percentiles. Accordingly, we show data averages for percentiles 78-100.

Figure B2 (a) plots the percentage change in real hourly wages between 1960 and 2017 in the ACS and Census against the model. The model captures both the negative slope of the low-skill households, with the flattening out of the curve towards the upper end of the distribution and the wage increase for high-skill households. Replicating the U-shaped polarization pattern, the bottom and top of the distribution experience an income growth 8% and 15% higher than the average, respectively, while large parts of the middle of the distribution experience relative income declines. Figure B2 (b) shows changes in the Törnqvist price index between 1960 and 2017 in the model and in the data. Be reminded that the model does not target the empirical differences in price indices directly, but matches the evolution of relative prices, the difference in expenditure patterns and the rise in income inequality. With these targets matched, the model generates a downward sloping pattern of the price index as in the data, although slightly less steep, indicating that the purchasing power of richer households has increased relative to poorer households.

Lastly, we focus on consumption inequality. In Figure B2 (c), we show expenditures that are normalized in the same way as panels (a) and (b). We use data from Jaravel and Lashkari (2024). In panel (d), we deflate in addition each percentile of consumption expenditures with the percentile-specific price index of panel (b). This means we measure changes in consumption expenditures by the price that people actually need to pay for their bundles. This approach approximates the

Figure B2: Data and model along the income distribution



Note: The figure compares key model variables (dots/dotted lines) with data moments (solid line/dashed line). Percentiles 1–77 show the distribution of low-skill workers, percentiles 78–100 the homogeneous high-skill workers. The data and model series are normalized by the low-skill mean. Panel (a) shows log changes in income, (b) the Törnqvist price index, (c) log expenditure increase and (d) log expenditure increase deflated with the percentile-specific price index. We use our CEX data from Section 2.5, where we calculated price changes between 1984 and 2017. Because the time frame is longer now, we multiply each value with a factor 57/33. Sources: ACS, Census, NIPA, CEX, Jaravel and Lashkari (2024).

real costs of living of the households, which is the more relevant concept for consumers.

We match well the data for the low-skill households. For high-skill households, the model captures that their expenditures are above the other percentiles, but the difference is slightly too large.

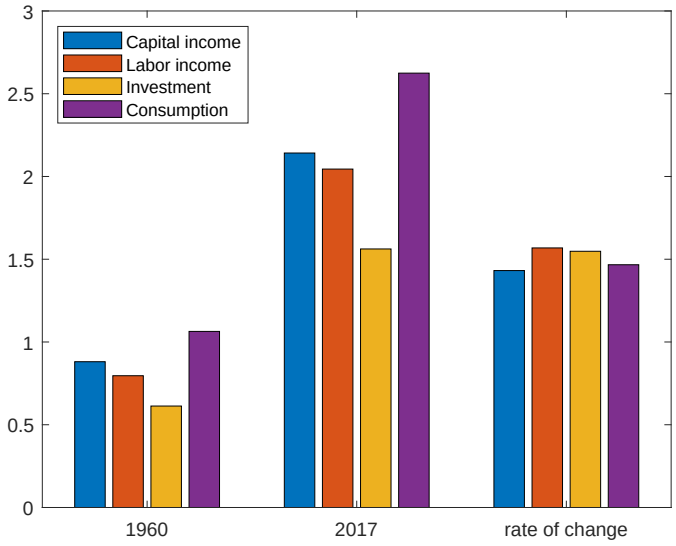
While expenditures in panel (c) differ by about 10 percent between the top and bottom of low-skill workers, this difference shrinks to almost zero when adjusting for percentile-specific price changes. In contrast, the difference between high-skill workers and the first percentile increases from about 10 percent to about 20 percent. While actual expenditures move closely with the evolution of income for lower income groups, the price effect partly overturns the pure income effect and leads to a different shape of the increase in real expenditures. This is especially true for the bottom part of the distribution that suffers a relative decline in purchasing power. This shows that it is important to consider real units when discussing consumption inequality. To the workers, it is not nominal expenditures that matter, but what they can actually buy with the money they spend.

B.5 The Role of Capital Income and Investment

According to eq.(19), high-skill households receive two types of income, labor and capital income, which they need to allocate across income and investment. Is this choice consequential? Figure B3 shows the different types of incomes and expenditures for 1960 and for 2017 and their percentage change. In both years, capital and labor income each make for about half of total income of high-skilled households. They both increase to about 2.5 times the initial value. So the increasing skill premium and the rising capital share move in tandem. Jaimovich et al. (2021) and Eden and Gaggl (2018) also stress the importance of the increase in the capital share, and our findings are in line with theirs. Given that changes in both the labor share and the skill premium were targeted in the model calibration, this is not surprising.

High-skill households allocate about 60% of their income to consumption and the rest to investment. Importantly, both consumption and investment grow by about the same rate between 1960 and 2017. As the model is solved with a focus on changes, the fact that all four variables grow by about the same rate implies that it does not make a great difference for the income effect whether capital and investment are included or excluded.

Figure B3: Income and expenditure of high-skill households



Note: The graph shows model predictions for per-capital capital income, labor income, investment and consumption of high-skill workers. Numbers are per-capita.

B.6 Counterfactual Exercise with Muted Labor Supply Growth

Table B2 compares model outcomes and data moments for a counterfactual where we mute labor supply growth. This means that both high-skill and low-skill labor stay at their 1960 value during the whole simulation period. This scenario results in a more pronounced income polarization. On the one hand, the skill premium now grows much faster. On the other hand, the premium of routine over manual wages declines much more. This result complements the exercise in Section 5.2, in which the labor demand increase was muted. In absence of an increase in (primarily high-skill) labor supply, the effects of digitalization on incomes would have been more extreme. Also noteworthy is that the labor share stays almost completely constant in this scenario. The declining labor share is thus not only driven by technological advances, but also by the fact that labor supply has grown.

Muting the increase in labor supply also results in a lower GDP per capita than in the baseline.

Table B2: Counterfactual $g_L = g_H = 0$, moments vs. data

Moment	Data	Model
Skill premium, 1960-2017	39.9	238.6
Relative price, 1960-2017	244.2	203.2
ICT intensity	643.1	384.3
Relative consumption, 1960-2017	-85.4	-91.8
Labor share, 1960-2017	-11.7	-2.7
Rate of return to ICT capital, 1960-2017	-97.8	-98.1
Median increase in exp. share good 2, 1960-2017	70.7	60.3
Percent diff., exp. share good 2 Top 1 Bot 1, 2017	23.3	28.1
Routine-manual wage premium, 1960-2017	-31.7	-49.1
GDP per capita, 1960-2017	281.6	200.4
Value-added share of sector 2, 1960-2017	71.4	52.3

Note: All moments are expressed in percentage changes, except for the difference in expenditure shares.

This shows that complementarities between technology and labor are of crucial importance for GDP growth.

Lastly, note that the muted growth in labor supply has almost no effect on the relative price and on consumption patterns. Labor supply changes do not explain any of the price effect .