Policy Research Working Paper 7487

On the Welfare Implications of Automation

Maya Eden Paul Gaggl



WORLD BANK GROUP

Development Research Group Macroeconomics and Growth Team November 2015

Policy Research Working Paper 7487

Abstract

This paper establishes that the rise in the income share of information and communication technology accounts for half of the decline in labor income share in the United States. This decline can be decomposed into a sharp decline in the income share of "routine" labor—which is relatively more prone to automation—and a milder rise in the non-routine share. Quantitatively, this decomposition suggests large effects of information and communication technology on the income distribution within labor, but only moderate effects on the distribution of income between capital and labor. A production structure calibrated to match these trends suggests modest aggregate welfare gains from automation.

This paper is a product of the Macroeconomics and Growth Team, Development Research Group. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at http://econ.worldbank.org. The authors may be contacted at meden@worldbank.org and pgaggl@uncc.edu.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

On the Welfare Implications of Automation

Maya Eden¹

The World Bank Development Research Group Macroeconomics and Growth Team 1818 H St. NW Washington, DC 20433 Email: meden@worldbank.org

Paul Gaggl¹

University of North Carolina at Charlotte Belk College of Business Department of Economics 9201 University City Blvd Charlotte, NC 28223-0001 Email: pgaggl@uncc.edu

JEL: E25, E22, J24, J31, O33 Keywords: labor share, capital share, job polarization, ICT

¹We are grateful to Eric Bartelsman, David Berger, Ben Bridgman, Bart Hobijn, Roberto Fattal Jaef, Nir Jaimovich, Aart Kraay, Richard Rogerson, Juan Rubio-Ramirez, Luis Serven, Michael Sposi, Alan M. Taylor as well as seminar participants at the World Bank, Vanderbilt University, Duke University, the Stockholm School of Economics (SITE), the ABCDE Conference on the Role of Theory in Development Economics, the 2014 NBER Summer Institute (CRIW), and the 2014 SEA Meetings for extremely helpful comments and suggestions.

1. Introduction

The increasing importance of computers in production has raised many questions regarding both the distributional and the aggregate implications of automation. Many have voiced concerns regarding potentially adverse effects of automation on the distribution of income across different types of labor, as well as on the distribution of income between capital and labor.¹ In parallel, there has been extensive debate regarding the contribution of information and communication technology (ICT) to output and productivity.² Overall, these issues are at the core of a broader open question: what are the overall net welfare gains or losses from the "ICT revolution"?

Our analysis contributes to the understanding of this matter in two ways. First, we quantify the extent to which the decline in the labor income share has been directly countered by an increase in the ICT capital income share. This measurement exercise suggests that about half of the decline in the labor income share can be attributed to automation. Interestingly, this estimate is consistent with the findings of Karabarbounis and Neiman (2014), who arrive at this conclusion using a very different methodology.³

Second, we use disaggregated trends in capital and labor income shares to calibrate an aggregate production function that emphasizes the interactions between ICT and different types of labor. Within a representative agent framework, this production structure suggests that the welfare gains from ICT are equivalent to a relatively modest 3.6% permanent increase in consumption from the perspective of 1968. Moreover, given potentially adverse distributional implications, this

¹For example, Karabarbounis and Neiman (2014), Elsby, Hobijn and Sahin (2013), and Bridgman (2014) study the implications of ICT for the distribution between capital and labor. On the other hand, Autor and Dorn (2013), Akerman, Gaarder and Mogstad (2013), and Gaggl and Wright (2015) are examples of recent studies that analyze the implications of ICT for the distribution of income within labor—across different types of labor. See Acemoglu and Autor (2011) for an extensive review of the broader literature surrounding this topic.

²See for example Colecchia and Schreyer (2002), Basu, Fernald, Oulton and Srinivasan (2003), Jorgenson and Vu (2007), Bloom, Sadun and Van Reenen (2012), as well as Acemoglu, Autor, Dorn, Hanson and Price (2014) and references therein.

³Karabarbounis and Neiman (2014) exploit cross country variation in the relative trends of the labor income share and the investment price to calibrate an elasticity of substitution between (aggregate) capital and (aggregate) labor. Their calibrated production function suggests that the declining investment price accounts for roughly half of the decline in the labor income share. Instead, our methodology consists of measuring the decline in the labor income share that is directly countered by an increase in the ICT capital income share.



Notes: Occupation specific income shares are based on CPS earnings data from the annual march supplement (1968 and after) and rescaled to match the aggregate income share in the Non-Farm Business Sector (BLS). The underlying earnings data are topcode adjusted using Piketty and Saez's (2003) updated estimates of the US income distribution (PS). Non-routine workers are those employed in "management, business, and financial operations occupations", "professional and related occupations", and "service occupations". Routine workers are those in "sales and related occupations", "office and administrative support occupations", "production occupations", "transportation and material moving occupations", "construction and extraction occupations", and "installation, maintenance, and repair occupations" (Acemoglu and Autor, 2011). For details see Section 2. The construction of capital-type specific income shares is described in Section 3. The underlying data are nominal gross capital stocks and depreciation rates, drawn from the BEA's detailed fixed asset accounts. The vertical dashed line in panel B marks the year 1968.

estimate represents an upper bound for the overall welfare gains from the "ICT revolution".⁴

We begin with a decomposition of U.S. income shares over the period 1968-2013 (Figure 1 and Table 1). Our analysis builds on the organizing framework of Autor, Levy and Murnane (2003) who highlight the distinction between "routine" and "non-routine" occupations, where routine occupations are jobs which are relatively more prone to automation.⁵ This decomposition confirms that the relatively steady but mild decline in the aggregate labor income share—recently documented on a global scale by Karabarbounis and Neiman (2014)—masks a much steeper decline in the income share of routine labor. However, the declining share of routine labor was

⁴It is also worth noting that our welfare gains do not take into account the costs associated with organizational change (Brynjolfsson and Hitt, 2003; Gaggl and Wright, 2015) and the accumulation of intangible capital (Bloom et al., 2012), shown to significantly affect on the efficiency of using ICT, and should thus be taken as an upper bound.

⁵Autor et al.'s (2003) seminal work has spurred a substantial body of literature documenting that ICT complements "non-routine" tasks—involving hard to automate, often inter-personal skills—and it replaces "routine" tasks—ones that follow exact, pre-specified procedures (Acemoglu and Autor, 2011). While the majority of this literature documents conditional correlations, two recent studies by Akerman et al. (2013) and Gaggl and Wright (2015) provide direct, causal evidence for this view.

			Lab	Labor Share		ital Share
	Labor Share	Capital Share	Routine	Non-Routine	ICT	Non-ICT
1968	63.4	36.6	38.6	24.8	1.3	35.3
2013	57.1	42.9	23.6	33.6	4.1	38.8
Percentage Po	oint Change sinc	e 1968				
1968-2013	-6.3	6.3	-15.0	8.7	2.8	3.5

Table 1: The Division of Income in the US

Notes: The table summarizes the long run trends in labor and cpaital shares as depicted in Figure 1. See the notes to Table 1 for details on the data construction.

almost entirely offset by an increase in the income share of non-routine labor, leading to a much milder decline in the aggregate labor income share. The direct counterpart to this decline was a 2.8pp increase in the ICT share and a 3.5pp increase in the non-ICT share. We find that the increase in the non-ICT share was primarily driven by the rising housing share in the post 2001 period (Figure 10), which is unlikely related to automation. Taken together, these trends suggest that automation accounts for slightly less than half of the overall decline in the labor income share.

Mindful of potential distributional costs, the second part of the paper aims to quantify an upper bound on the benefits from automation. To this end, we use our disaggregated trends in income shares to calibrate an aggregate production function, which we embed in an otherwise standard neoclassical growth model. We quantify the impact of the ICT revolution by contrasting our baseline simulation—in which we assume the observed decline in ICT prices—with a counterfactual in which the price of ICT capital is held constant at its 1968 level.⁶

Our model suggests that, compared to the counterfactual, the declining ICT price increased steady state output by 12% and steady state consumption by 9.4%. These numbers are substantially lower than those suggested by Karabarbounis and Neiman (2014), who consider a production framework with homogeneous capital and labor inputs and estimate steady state output and

⁶Since it is well documented that the declining price of ICT is reflective of efficiency gains in the production of computing power, we interpret the relative price of ICT as a measure for this form of technological progress. For empirical evidence supporting this view, see for example Figure 1.1 in Hennessy, Patterson and Asanović (2012) as well as references therein. We provide a theoretical argument for this view in Section 5.

consumption gains of 22.8% and 20.1%, respectively.⁷ While Karabarbounis and Neiman (2014) focus on comparing welfare across steady states, we calculate welfare gains from the perspective of 1968, taking into account the transitional costs associated with accumulating higher steady state capital levels. Our analysis suggests welfare gains equivalent to a 3.6% permanent increase in consumption from the perspective of 1968. Comparable estimates using the production function in Karabarbounis and Neiman (2014) imply welfare gains of 5%.

In sum, our analysis suggests that both the costs and the benefits associated with the declining price of computing power may have been exaggerated. In terms of costs, the trends in income shares indicate that the redistribution of income associated with automation was mainly within labor, rather than between capital and labor. In this context, it is useful to observe that routine occupations consist mainly of middle income jobs, while non-routine occupations consist both of high-skilled professional jobs and low-skilled service jobs (Acemoglu and Autor, 2011). Thus, the welfare implications of this redistribution are ambiguous. From a distributional perspective, the decline in the aggregate labor income share is more worrisome; however, over half of it has been due to a rise in the non-ICT capital income share, which is less likely symptomatic of automation. In terms of benefits, our quantitative analysis suggests relatively modest welfare gains associated with automation, even from the viewpoint of a representative agent model in which redistribution is frictionless.

Our general equilibrium framework is closely related to Krusell, Ohanian, Ríos-Rull and Violante (2000), who study the implications of declining investment prices for the rising skill premium in the second half of the 20th century.⁸ We focus on the distribution of income across different types of tasks, rather than skills, following the more nuanced tasks framework first suggested

⁷The results are not entirely comparable with Karabarbounis and Neiman (2014), since they consider the global decline in the labor income share since 1980 while we focus on the decline in the labor income share in the US since 1968. In Section 5 we derive directly comparable numbers using their calibrated production function, which are 20% for output and 17% for consumption.

⁸See Goldin and Katz (1998, 2008) for a review of a broad literature analyzing the rising skill premium over that period.

by Autor et al. (2003).⁹ Within this organizing framework, routine and non-routine labor are effectively defined in terms of their substitutability with ICT. We therefore distinguish between ICT and non-ICT capital, rather than equipment and structures as in Krusell et al. (2000).¹⁰

The rest of the paper is organized as follows: Sections 2 and 3 describe the details of the measurement of the disaggregated trends in income shares displayed in Figure 1. These sections further offer a decomposition of the income shares into price and quantity components as well as a discussion of the role of international trade, housing, and industrial composition. Section 4 illustrates our calibration exercise and Section 5 describes our main quantitative analysis. Finally, Section 6 offers some concluding remarks.

2. Decomposing the Labor Income Share

To measure routine and non-routine labor income shares, we consider two alternative measures of earnings at the occupation level in the U.S. Current Population Survey (CPS): annual earnings from the March supplements (MARCH) starting in 1968 (provided by IPUMS, Ruggles, Alexander, Genadek, Goeken, Schroeder and Sobek, 2010), and weekly earnings for the outgoing rotation groups (MORG) starting in 1979 (provided by the NBER). Based on these earnings data we decompose the U.S. aggregate labor share (BLS) into the portion going to routine and non-routine labor, respectively.

⁹They document that the set of tasks that are most readily prone to automation are primarily performed by middleskill jobs, and these jobs have seen a relative decline in both employment and wages. In contrast, both high skill occupations—that require creative thinking—*and* low skill service jobs—that often require interpersonal skills—have gained along both dimensions. These facts are now well documented for many developed nations, starting with the work by Acemoglu (1999) for the US over the period 1983–1993, and in more recent periods by Goos and Manning (2007) for the UK, Goos, Manning and Salomons (2009) for 16 EU countries, and Autor, Katz and Kearney (2008) as well as Autor and Dorn (2013) for the US.

¹⁰While ICT capital is mostly included in equipment, the equipment category considered in Krusell et al. (2000) is much broader and includes vehicles, electrical machinery, and even furniture. Like Krusell et al. (2000), vom Lehn (2015) also studies the declining price of equipment. While his analysis of the labor market is also guided by the tasks framework (Autor et al., 2003), his approach differs from ours both in terms of the classification of occupations and the calibration approach. The most important distinction, however, is his focus on equipment, rather than ICT.

We note that this is not a trivial distinction for several reasons: first, the equipment share of output is roughly 25% (based on the calculations by Krusell et al., 2000), yet we estimate an ICT income share of at most 4% in recent years. Second, we further document that ICT accounts for essentially all of the decline in the relative price of equipment and that the price of non-ICT capital, a substantial part of which comprises equipment, was virtually constant relative to the GDP deflator throughout our entire sample (see panel A of Figure 8). Finally, we find that the income share of non-ICT is trend-less, yet that of ICT displays a strong upward trend (Figure 1).



Notes: Panel A contrasts aggregate income shares (as a fraction of GDP) based on the CPS outgoing rotation groups (MORG), aggregates reported in the NIPA tables, as well as a BLS estimate for the total non-farm business sector that includes benefits, self employed, proprietors income, and other non-salary labor income. The aggregate series are drawn from FRED. The series labeled "CPS (R+NR)" is constructed from our occupation specific earnings based on the monthly CPS MORG extracts provided by the NBER. The data are seasonally adjusted with the U.S. Census X11 method. Panel B contrasts the raw earnings reflected by CPS topcoded values and our series that adjust top-coded earnings with the appropriate (updated) estimates by Piketty and Saez (2003).

To do so, we use the CPS sampling weights and construct an estimate of the aggregate wage bill at the detailed occupation level, which requires several non-trivial adjustments to the raw data. First, since the U.S. Department of Labor's (DOL) classification of occupations changes several times during our sample period, we aggregate individuals into a panel of 330 consistent occupations, designed by Dorn (2009).¹¹ Second, and more crucial for our analysis, we follow Champagne and Kurmann (2012) and adjust top coded earnings based on Piketty and Saez's (2003) updated estimates of the cross-sectional income distribution.

Based on these adjusted earnings numbers, we then compute the aggregate annual wage bill and divide it by nominal GDP, to construct the share of wage and salary earnings in aggregate income. As illustrated in panel A of Figure 2, the aggregate labor share based on earnings data in the CPS-MORG accounts for stable 70% of the one based on total non-farm business labor income (which includes benefits, pensions, self employed income, etc.). Moreover, the two series are almost perfectly correlated over time.

¹¹We thank Nir Jaimovich for providing a crosswalk between Dorn's (2009) occupation codes and the latest Census classification that is used in the CPS since 2011. This crosswalk is the same as in Cortes, Jaimovich, Nekarda and Siu (2014).

To compute the routine and non-routine income shares we define routine and non-routine workers as suggested by Acemoglu and Autor (2011). That is, we consider workers employed in "management, business, and financial operations occupations", "professional and related occupations", and "service occupations" as non-routine; and we define routine workers as ones employed in "sales and related occupations", "office and administrative support occupations", "production occupations", "transportation and material moving occupations", "construction and extraction occupations", and "installation, maintenance, and repair occupations". Note that this classification emerges out of an extensive literature, surveyed in Acemoglu and Autor (2011), that originated from the seminal work by Autor et al. (2003). They and many other contributions in this line of research use detailed information on the task content of at least 300 detailed occupations (depending on the study) obtained from the Dictionary of Occupational Titles (DOT) and its successor O*Net. The classification used here is the "consensus aggregation" suggested by Acemoglu and Autor (2011) that captures the key insights from the more detailed micro analyses. We drop farm workers for all our analyses for comparability with the BLS measure of the labor income share. To compute the income shares corresponding to each occupation group, we simply compute the aggregate annual wage bill within each occupation group and divide it by nominal GDP.

Finally, we proportionately rescale both group specific income shares (which originally add up to the series labeled "CPS(R+NR)" in panel A of Figure 2) so that they match the share of (non-farm business) labor income in GDP, as estimated by the BLS (top line in panel A of Figure 2). The resulting routine and non-routine income shares using the MARCH earnings measure are displayed in panel A of Figure 1 and Figure B.13 in Appendix B illustrates that income shares based on the MORG earnings measure reveal a virtually identical picture.

This decomposition highlights a striking feature: in the U.S., the decline in the aggregate labor share is entirely accounted for by routine occupations, while the income share of non-routine labor has been rising.



Notes: Panel A plots the unconditional mean of log real annual earnings in each occupation group. Panel B graphs the coefficients from annual regressions of individual level log real earnings on a non-routine dummy and a host of demographic control variables, including flexible functional forms in industry, age, and education. Occupation and individual specific earnings are based on the the annual march supplements in the CPS (MARCH) provided by IPUMS (Ruggles et al., 2010). We deflate eranings data with the chain type implicit price deflator for personal consumption expenditures. Non-routine workers are those employed in "management, business, and financial operations occupations", "professional and related occupations", and "service occupations". Routine workers are those in "sales and related occupations", "office and administrative support occupations", "production occupations", "transportation and material moving occupations", "construction and extraction occupations", and "installation, maintenance, and repair occupations" (Acemoglu and Autor, 2011).

2.1. Price-Quantity Decomposition

The declining routine labor income share relative to the non-routine labor income share could be driven either by a change in relative wages, a change in relative labor inputs, or both. We find that both an increasing non-routine wage premium and an increase in non-routine labor inputs contribute to this trend.

A large body of literature has documented increasing wage *polarization* over the past three decades—a relative increase in wages for high- and low-paying jobs relative to middle-income jobs.¹² Closely related, we provide evidence of an increasing non-routine wage premium, providing further support for this view.

As a baseline reference, we start with estimating simple averages of log real earnings for each type of labor. Panel A of Figure 3 illustrates these estimates and gives a first indication of a steadily increasing wedge between non-routine and routine pay. However, to ensure that this wedge is not

¹²See Acemoglu and Autor (2011) for a comprehensive summary of this literature.

simply driven by a changing composition of characteristics of routine and non-routine workers, we estimate the following set of cross-sectional wage regressions separately for each year, *t*:

$$\ln w_{i,t} = \beta_{0,t} + \beta_{1,t} N R_{i,t} + \beta_{2,t} X_{i,t} + \epsilon_{i,t} \quad \text{for } t \in \{1968, \dots, 2013\},\tag{1}$$

where $NR_{i,t}$ is a dummy variable indicating that individual *i* works a non-routine job in year *t* and $X_{i,t}$ includes a variety of control variables. In particular, we include gender, race, and full time employment dummies, we control for the weeks worked, a full set of industry fixed effects (50 industries constructed from SIC industry codes by the NBER), as well as fourth order polynomials in age, education, and the interaction of education and age.

We estimate regressions (1) based on individual level data from the annual CPS march supplements and weight by the CPS sampling weights.¹³ Panel B of Figure 3 plots the resulting time series of estimates $\hat{\beta}_{1,q}$ and the associated 95% confidence intervals based on standard errors that are clustered on industry. These estimates highlight that the rising relative wage for non-routine labor is not entirely driven by the rising skill premium or by specific industries. The latter observation is of particular importance, as it highlights that the wage premium is not simply due the steady decline in manufacturing as the estimates $\hat{\beta}_{1,q}$ are identified from within industry variation. These estimates therefore suggest that part of the increase in the non-routine income share is driven by a steadily increasing gap between routine and non-routine pay.

Our estimates suggest that the non-routine wage premium rose from around 14% in 1968 to more than 26% in 2013. Given the sheer size of this premium we find it unlikely to be driven by frictions (like barriers to entry). In our calibration exercise, we will therefore assume flexible labor markets and attribute the non-routine wage premium to some form of unobserved ability (e.g. managerial skills, "people" skills, etc.).

Similarly, to measure the trends in labor inputs we distinguish between raw labor (employment) and effective units of labor which take into account differences in worker attributes. In a

¹³Note that in an earlier version we estimated these regressions at the quarterly level based on the CPS MORG. The results are qualitatively equivalent but we prefer the longer time horizon provided by the annual march supplements. The CPS MORG results are available form the authors upon request.



Notes: Panel A plots employment levels in routine and non-routine jobs as reflected in the CPS. The graph plots both the raw CPS numbers as well as our imputed "effective" units based on equations (4) and (5). Panel B illustrates the relative importance of non-routine jobs.

frictionless world, an "effective unit" of labor needs to be paid a fixed wage w_t and therefore, the ratio of non-routine to routine labor in effective units is given by

$$\frac{s_{r,t}}{s_{ln,t}} = \frac{w_{r,t}L_{r,t}}{w_{nr,t}L_{nr,t}} = \frac{w_t(e_{r,t}L_{r,t})}{w_t(e_{nr,t}L_{nr,t})} = \frac{L_{r,t}^e}{L_{nr,t}^e},$$
(2)

where $e_{r,t}$ and $e_{nr,t}$ are effective units of labor embodied in routine and non-routine workers, respectively. Panel B of Figure 4 illustrates the time path of this ratio, revealing that non-routine labor inputs have increased relative to routine labor inputs, both in terms of raw labor and in terms of effective units of labor.

For our calibration exercise, it will be useful to construct the time paths of routine and nonroutine labor in effective units, as well as a sequence of real wages per effective unit of labor. We obtain an estimate for w_t directly from regression model (1). To do so, we normalize X_{it} such that $X_{it} = 0$ corresponds to our "reference worker": a 19-year-old, white, male, full-time, routine, manufacturing worker with a high school degree, who works 50-52 weeks per year. This allows us to interpret our estimated constant in regression (1) as the baseline wage, that is $\hat{w}_t = \hat{\beta}_{0,t}$. Figure 5 illustrates the resulting baseline real wage series and we note that our estimates are in line with a vast literature in labor economics that finds stagnating real wages for the "middle class" starting



Notes: The figure plots $\hat{w}_t = \hat{\beta}_{0,t}$ based on estimates from regression model (1) with $X_{it} = 0$ for our baseline worker: a 19-year-old, white, male, full-time, routine, manufacturing worker with a high school degree, who works 50-52 weeks per year. Real earnings are based on the U.S. GDP deflator with 1979=1.

in the early 1970s (Levy and Murnane, 1992).

Based on the identity $w_t L_t^e = \sum_i w_{it} L_{it}$, aggregate effective employment is then given by

$$L_t^e = \frac{\sum_i w_{it} L_{it}}{w_t} \tag{3}$$

which allows us to compute the effective levels of routine and non-routine employment as

$$L^{e}_{nr,t} = \frac{1}{1 + \frac{s_{r,t}}{s_{ln,t}}} L^{e}_{t}$$
(4)

$$L_{r,t}^e = L_t^e - L_{nr,t}^e \tag{5}$$

Panel A of Figure 4 illustrates the time paths for routine and non-routine employment, both in actual and effective units of labor.

Figures 3 and 4 clearly illustrate that the increase in non-routine labor's share in income is due to a substantial increase in both the non-routine wage premium as well as non-routine employment relative to routine employment.

2.2. Potential Drivers of the Trends in Labor Shares

For our quantitative analysis in Sections 4 and 5, we will adopt the view that the differential trends in routine and non-routine labor income shares are driven by different interactions with ICT. Of course, there are other potential causes for these trends. While far from conclusive, this section discusses three plausible alternative explanations and offers some evidence suggesting that these are not likely the dominant forces behind the observed trends displayed in panel A of Figure 1.

First, at least since the seminal work of Piketty and Saez (2003), a growing literature and heated public debate have put much emphasis on the rising gap between the earnings of the top 1% of earners relative to the remaining 99% (e.g., Jones and Kim, 2014; Atkinson, Piketty and Saez, 2011). To gauge at the potential impact of this recent divergence on our reported trends, Panel B of Figure 2 illustrates the impact of our top-code adjustment relative to the raw values reported in the CPS.¹⁴ It is important to note that, despite the fact that we adjust top-coded values at the individual level, our adjustment effectively results in a level shift in aggregate income shares; all the original trends are preserved after the adjustment. If the divergence of the top 1% were really driving the differential trends for routine and non-routine income shares, then we should expect a strongly rising gap between the adjusted and raw numbers for non-routine workers but not for routine workers. As this is not readily apparent in panel B of Figure 2 we argue that "the one percent" are not likely the main driver of our results. Nevertheless, panel B of Figure 2 highlights that the top-code adjustment is crucial in order to construct quantitatively meaningful aggregate income shares based on earnings reported in the CPS.

Second, it also appears plausible that the decline in routine labor's importance in production might merely be a symptom of the decline in manufacturing jobs. However, Figure 6 illustrates that the strong divergence between routine and non-routine income shares is not only due to the disappearance of classic blue collar jobs—which are primarily concentrated in the manufacturing sector. Specifically, at least since 1988, both service and managerial/professional occupations gained about 20% in income share. Thus, both abstract and manual occupations contributed about equally to the rise in the non-routine income share since 1988. Similarly, while traditional blue col-

¹⁴Note that these shares are not re-scaled to match the BLS's estimate of the aggregate labor share.



Notes: Panel A plots unadjusted income shares of four major occupation groups as reflected in the CPS MORG. These shares do not add up to 100% as earnings reflected in the CPS MORG do not properly measure benefits as well as proprietors' income. Panel B imposes the normalization 1988q1=1 to illustrate the clear divide between routine and non-routine tasks. The dashed vertical line indicates 1988q1.

lar occupations lost about 40%, administrative/clerical occupations lost a little more than 20% in their income share. This makes clear that, while traditional blue collar occupations are by far the dominant driver of the decline in the routine labor income share, administrative/clerical jobs still explain at least a quarter of the decline in the routine share since 1988—clearly a non-negligible portion. Moreover, our estimates of the non-routine wage premium based on equation (1) control for industry fixed effects, suggesting that the premium for non-routine work captures a withinindustry feature.

Third, it is possible that international trade has played a key role in shaping the allocation of labor across routine and non-routine occupations (Autor, Dorn and Hanson, 2013; Elsby et al., 2013). Specifically, if most of the output produced by non-routine labor in recent years is in fact exported—such as management and other business consulting services—and if routine-intensive products are now primarily imported, then the true driver of our measured income shares would in fact be international trade and not necessarily ICT.

To evaluate this hypothesis we calculate the share of labor expenditures that is embodied in (net) exported goods and services. Define $s_{R,i,t}$ and $s_{NR,i,t}$ as the routine and non-routine labor shares in industry *i*, respectively. Further define $NX_{i,t}$ the total expenditure on net exports in



Figure 7: The Role of Trade for Labor Shares

Notes: The Figure illustrates the routine and non-routine labor shares, adjusted for the role of exports. Note that the labor share is not rescaled to match the official BLS labor share and only incorporates information on earning from the CPS MORG.

industry *i*. Thus, we can define the share of expenditure on routine and non-routine labor, $\ell \in \{R, NR\}$, embodied in net exports as follows:

$$s_{NX,\ell,t} = \frac{\sum_{i \in I} s_{\ell,i,t} \cdot NX_{i,t}}{GDP_t} \tag{6}$$

To compute the industry specific labor income shares we use earnings data from the CPS MORG, and value added data by industry from the BEA. Trade in goods and services by industry is also obtained from the BEA.

Based on these calculations, we illustrate in Appendix C that the US is a net exporter of labor services. However, this amounts to only a small fraction of the total labor expenditure in the US. Figure 7 illustrates this finding, highlighting that trends in the net exports of goods and services only explain a very small fraction of the trends in the aggregate routine and non-routine labor shares.

3. Decomposing the Capital Income Share

While we were able to measure labor income shares directly from observed earnings data, estimating the payment to different types of capital requires a more structural approach. With a single type of physical capital—as is customary in most macro analyses—it is straightforward to measure the capital income share as one minus the labor income share. However, we are interested in separating the payments to several distinct types of capital. The payment to each type of capital is comprised of both a unit payment (the rental rate of capital) and the (real) stock of capital analogous to the wage rate and the physical amount of labor provided by the worker. Both of these items are challenging to measure, especially when the relative prices of the various types of capital are changing over time.

We build on two standard assumptions, that have been used to measure the returns to capital at least since the seminal work by Hall and Jorgenson (1967) and Christensen and Jorgenson (1969), which allow us to directly measure capital type specific income shares from the BEA's current cost values for the stocks of detailed assets in the US.¹⁵

Specifically, we impose a no-arbitrage condition in investment in addition to the standard constant returns to scale assumption.¹⁶ Suppose that there are several types of capital, denoted K_i , with i = 1, ..., I. If the production technology exhibits constant returns to scale in all factors, the share of payments to capital must satisfy the following equilibrium relation:

$$s_{K,t} = \sum_{i}^{I} \frac{R_{i,t} K_{i,t}}{P_t Y_t} = 1 - s_{L,t},$$
(7)

where $s_{K,t}$ and $s_{L,t}$ denote the aggregate capital and labor income shares, respectively, Y_t is final

¹⁵Official documenation for the BEA's methodology to construct these estimates is available at http://www.bea.gov/national/pdf/Fixed_Assets_1925_97.pdf. Most macroecnomic studies using capital stocks utilize a simpler version of the perpetual inventory method than the BEA's estimates, usually based on linear constant depreciation and aggregate real investment rates. We prefer the BEA's estimates for several reasons: first, they are provided at the detailed asset level; second, they allow for time varying non-linear depreciation patterns; finally, these estimates allow us to directly use nominal stocks at current cost, rather than chain-weighted quantity indexes.

¹⁶For a few more recent contributions that use the same basic strategy to compute the return to specific types of capital in various contexts see for example Jorgenson (1995), O'Mahony and Van Ark (2003), and Caselli and Feyrer (2007). We outline the basic idea of our implementation to measure capital type specific income shares here and provide detailed derivations in Appendix A.

output with associated price P_t , and $R_{i,t}$ denotes the (nominal) rental rate of capital type *i*. Note that, if factor markets are competitive, the real rental rate $(R_{i,t}/P_t)$ is equal to the physical marginal product of a unit of capital. However, our approach does not require this assumption.

Suppose an investor bought a unit of capital type *i* at price $P_{i,t-1}$ and rents it out for a period at the real rental rate $R_{i,t}$. The gross return after production and re-sale of this piece of capital is then given by

$$\frac{R_{i,t} + P_{i,t}(1 - \delta_{i,t})}{P_{i,t-1}},$$
(8)

where $\delta_{i,t}$ is the depreciation rate for capital type i.¹⁷ In an equilibrium in which investors can choose between different assets, the return on each type of capital must equal the prevailing gross return on investment.

It is important to note that this does not require the (physical) marginal product of each type of capital to be equalized. In the standard neoclassical growth model, marginal products need to equalize since capital has a constant price relative to output and there is only one rate of depreciation. In our context, both the price as well as the depreciation rate of ICT is changing drastically relative to output and all other forms of capital (see Figure 8). Thus, no-arbitrage in investment requires the *gross return* (8) to be equalized across all types of capital.

We show in Appendix A how equations (7) and (8) allow us to compute the income share for each type of capital, defined as $s_{i,t} = \frac{R_{i,t}K_{i,t}}{P_t Y_t}$, based on the labor income share, nominal current cost values for each type of capital, capital specific depreciation rates, $\delta_{i,t}$, and a price index for ICT capital. To measure the current cost values of different types of assets we use the BEA's detailed fixed asset accounts and aggregate the BEA's detailed industry level estimates into three types of capital:¹⁵ we distinguish residential and non-residential assets as well as consumer durables according to the BEA's definition. Within the non-residential category and consumer durables we separate ICT and non-ICT assets. Specifically, within non-residential assets we consider an asset to be ICT if the BEA classifies it as software (classification codes starting with RD2 and RD4) or as equipment related to computers (classifications codes starting with EP and EN). See

¹⁷Notice that we use the same timing as Caselli and Feyrer (2007) here.



Notes: Panel A graphs implicit price deflators by capital type expressed as a fraction of the GDP deflator, which were constructed directly from the BEA's detailed fixed-asst accounts. The BEA GDP deflator is taken from FRED. Panel B depicts asset-specific depreciation rates constructed directly from the BEA's fixed asset accounts. See Tables F.7 and F.8 in Appendix A and the text for our grouping of assets. The dashed vertical lines indicate the year 1968.

Tables F.7 and F.8 in Appendix A for complete lists of the detailed assets grouped into the two types of non-residential capital. Within consumer durables we classify the following assets as ICT: PCs and peripherals (1RGPC); software and accessories (1RGCS); calculators, typewriters, other information equipment (1RGCA); telephone and fax machines (1OD50).

To construct the income shares of these assets, we further need an estimate of both the depreciation rate, $\delta_{i,t}$, and expected capital gains, $E[P_{i,t+1}/P_{i,t}]$, for each type of capital. We measure depreciation rates directly form the BEA's nominal values of depreciation for each type of detailed asset.¹⁸ We then employ implicit price deflators that we construct for each type of capital based on chain type price indices provided by the BEA, to measure capital type specific inflation.¹⁹ Panel A of Figure 8 depicts the path of prices for different types of assets, where we have aggregated residential and non-residential non-ICT capital, since the prices for these types of assets largely

¹⁸In particular, we measure depreciation rates based on the BEA's nominal values for depreciation and net capital stocks. That is, we compute $\delta_{i,t} = (P_{i,t}Dep_{i,t})/(P_{i,t}(NetStock_{i,t} + Dep_{i,t}))$. Since both measures are reported in year-end nominal values, the price terms cancel.

¹⁹Notice that this involves constructing appropriate chain type quantity aggregates and associated implicit price deflators for each capital type, derived from the BEA's estimates of stocks and prices for the detailed assets listed in Tables F.7 and F.8 in Appendix A.

evolve in lockstep.²⁰ This figure reveals two striking insights: First, the relative price of non-ICT capital and output/consumption are essentially constant throughout the entire sample. Second, the price of ICT capital falls substantially, both in absolute terms and relative to all remaining types of assets, especially after the 1982 recession. Panel B of Figure 8 graphs the respective depreciation rates for each type of capital.

Based on these measures, we use the derivations in Appendix A to construct the income share for each type of capital and Figure 1 illustrates the resulting estimates. One can clearly see that the income share of non-ICT capital does not show any significant trend throughout the entire sample. During the same period, the income share of ICT capital has roughly quadrupled, from around 1% in 1968 to around 4% in 2000. This suggests that the introduction of ICT did not significantly crowd out other forms of capital, whose income share fluctuated around a trend-less long run average of around 34%; instead, it crowded out routine labor.

3.1. Price-Quantity Decomposition

The above results suggest a substantial increase in the income share of ICT capital relative to the income share of non-ICT capital. While we did not use the real stock of capital to compute its income share, $s_{i,t}$, we find it instructive to decompose the payments to capital, $R_{i,t}K_{i,t}$, into a price and quantity component. To this end, we construct a chained quantity index for the stock of both ICT and non-ICT capital based on the BEA's fixed asset accounts. Panel A of Figure 9 shows that the stock of ICT capital in 2012 is about 40 times its 1968 level. On the other hand, Panel B illustrates that the rental rate of ICT capital—measured in units of final output—fell substantially over the same period. In particular, while an additional unit of ICT capital increased final output by about 30% in 1968 it only produced less than an additional 10% in 2012. This price-quantity decomposition suggests that the increase in the income share of ICT capital is due to massive accumulation of ICT, while the rental rate of ICT capital fell during this time.

²⁰See Figure D.15 in Appendix D for a decomposition of non-ICT assets into residential and non-residential assets.



Notes: Panel A graphs the stock of ICT and non-ICT captial relative to its 1968 level. Panel B depicts asset-specific real rental rates $(R_{i,t}/P_t)$ in % of final output, derived from expressions (A.5)-(A.6) in Appendix A. The underyling data are the BEA's detailed fixed-asset accounts. The dashed vertical lines indicate the year 1968.

3.2. The Role of Housing

Rognlie (2015) has recently argued that housing may be the main driver for an increase in the net capital share—the capital income share, net of depreciation—since 1970. In light of this, we briefly discuss the role of housing within the context of our analysis. To this end, we use the methodology described in Section 3 to decompose the non-ICT share into residential and non-residential assets. Figure 10 illustrates the resulting decomposition, again based on the BEA's fixed asset accounts.²¹

This figure highlights that the non-residential non-ICT capital share is completely trend-less throughout the period. Consistent with the findings of Rognlie (2015), the increase in the non-ICT capital income share in the post 2001 period is accounted for entirely by residential capital. Since the rising residential capital share is unlikely to be reflective of automation, our calibration assumes that the decline in the labor share which is attributable to ICT accumulation is only the portion that is directly countered by an increase in the ICT share.

²¹Note that this more disaggregated decomposition takes into account heterogeneous prices and depreciation rates for residential and non-residenital assets as measured by the BEA and depicted in Figure D.15 in Appendix D.



Figure 10: Capital Income Share: ICT/Residential/Non-Residential

Notes: The figures displays a decomposition of the capital share with three types of assets: ICT, residential, and non-residential non-ICT. The computations are based on the methodology described in Section 3. The dashed vertical line indicates the year 2001.

3.3. The Role of Industrial Composition

In Section 4, we will calibrate a one-sector aggregate production structure that matches the trends depicted in Figure 1. To justify such an approach, we illustrate that these trends are not merely a result of changes in the industrial composition but rather a within-industry phenomenon.

While a number of contributions in the broad literature surveyed by Acemoglu and Autor (2011) have convincingly shown that differential trends in earnings and employment of routine and non-routine occupations are to a large extent driven by within industry variation, there is no evidence that the same is true for ICT and non-ICT shares. To address this question, we construct industry specific ICT and non-ICT shares as in Eden and Gaggl (2015) based on the BEA's detailed fixed asset accounts. We then estimate the average annual growth rates in these income shares both with and without industry fixed effects.

Table 2 illustrates that the ICT share was growing at about 3 % annually while the non-ICT share showed no significant trend growth over the period 1968 - 2012. Importantly, columns

	ICT S	Share	Non-IC	Non-ICT Share		
	(1)	(2)	(3)	(4)		
Ann. Trend Growth (%)	3.197*** (0.444)	3.197*** (0.233)	0.0534 (0.319)	0.0534 (0.0700)		
Industry FEs Observations	855	yes 855	855	yes 855		

Table 2: Trends in Capital Income Shares

Notes: The table shows regressions of 100 times the log ICT and non-ICT shares on a time trend for an annual panel of 19 broad sectors over the period 1968-2012. Columns 1 and 3 report pooled regressions while columns 2 and 4 condition on a complete set of fixed effects for 19 broad sectors. Sector specific income shares are based on the BEA fixed asset accounts and the BEA's estimates of sector specific value added and constructed as in Eden and Gaggl (2015). Standard errors are HAC robust and significance levels are indicated by * p < 0.1, ** p < 0.05, and *** p < 0.01.

3 and 4 illustrate that this finding does not change after we control for a complete set of fixed effects. In fact, we find that the levels of ICT intensity vary vastly across the various industries, but the differential trends in ICT and non-ICT income shares are predominantly a within industry phenomenon.

4. A Calibrated Production Function

This section lays out a procedure for calibrating an aggregate production function based on observed trends in income shares. The underlying assumption is that long-run trends in the *relative* income shares of ICT capital, routine- and non-routine labor are driven by the interactions between labor and ICT capital.

Using lower case letters to indicate variables expressed in per-worker terms, we assume an aggregate production function given by:

$$y = k_n^{\alpha} x^{1-\alpha},\tag{9}$$

where y is output, k_n is non-ICT capital and x is an input produced by routine inputs (x_r) and

non-routine inputs (x_{nr}) according to a constant elasticity of substitution (CES) aggregate:

$$x = (\eta x_r^{\theta} + (1 - \eta) x_{nr}^{\theta})^{\frac{1}{\theta}},$$
(10)

where $\eta \in [0, 1]$ and $\theta \leq 1$. Routine and non-routine inputs (x_i) are CES aggregates of ICT capital (denoted $k_{c,i}$) and labor (denoted l_i):

$$x_{i} = (\gamma_{i}k_{c,i}^{\sigma_{i}} + (1 - \gamma_{i})l_{i}^{\sigma_{i}})^{\frac{1}{\sigma_{i}}}).$$
(11)

where $\gamma_i \in [0, 1]$ and $\sigma_i \leq 1$. Note that the variable $k_{c,i}$ denotes the ICT capital inputs employed in the production of x_i . Using this notation, the aggregate supply of ICT capital is given by $k_c = k_{c,r} + k_{c,nr}$.

This specification allows for ICT to interact directly both with routine and non-routine labor inputs. As special cases, this double-nested-CES production function embeds the single-nested-CES specifications previously considered in the literature. For example, Autor and Dorn (2013) consider the following nested CES specification:

$$x = \left(\eta (\gamma_r k_c^{\sigma_r} + (1 - \gamma_r) l_r^{\sigma_r})^{\frac{\theta}{\sigma_r}} + (1 - \eta) l_{nr}^{\theta}\right)^{\frac{1}{\theta}}$$
(12)

The interpretation is one in which non-routine labor interacts with "routine inputs", which can be produced by either routine labor or ICT capital. For example, a (non-routine) applied economist is more productive if there is more data available; it is immaterial whether the data is collected by (routine) human surveyors or by an online survey. The complementarity between non-routine labor and ICT happens only through increasing the supply of routine inputs, which are complementary to non-routine labor. Routine labor is directly substitutable with ICT.

An alternative specification, in the spirit of Krusell et al. (2000), is a nested CES in which nonroutine labor interacts directly with ICT:

$$x = (\eta l_r^{\theta} + (1 - \eta)(\gamma_{nr} k_c^{\sigma_{nr}} + (1 - \gamma_{nr}) l_{nr}^{\sigma_{nr}})^{\frac{\theta}{\sigma_r}})^{\frac{1}{\theta}}$$
(13)

Here, non-routine labor is directly complementary to ICT capital, and routine labor is substitutable with a "non-routine" input that is produced jointly by ICT capital and non-routine labor. To set ideas, consider the case of word processing. Secretaries producing printed documents using type-writers are substitutable with writers directly producing their documents using word processors. In this case, the direct interaction between non-routine labor and ICT capital makes routine labor redundant. More generally, this functional form captures technologies that directly complement non-routine labor, performing tasks that would be less feasible for routine labor. For example, complicated computations such as a multivariate regressions would require a lot of patience and brainpower without a computer, rendering them practically unfeasible for routine labor. Similarly, communication technologies that allow for instantaneous sharing of ideas and information provide a service that is not feasibly provided by human messengers.

Both specifications are plausible and capture different interactions between labor and ICT. Both allow for substitutability between routine labor and ICT, and complementarity between nonroutine labor and ICT. We therefore remain a-priori agnostic and consider a flexible double-nested-CES specification. Note that equations 12 and 13 are embedded in our specification with $\gamma_{nr} = 0$ and $\gamma_r = 0$, respectively.

Given a set of parameters and aggregate quantities of labor and ICT capital, we can compute the equilibrium allocation of inputs under the assumption that marginal products are equalized across the production of x_r and the production of x_{nr} . Given the allocation of inputs, we can compute the implied routine labor income share, the non-routine labor income share and the ICT capital income share, all relative to the aggregate expenditure on x.²²

Note that our calibration strategy assumes that wages are equalized across routine and nonroutine occupations. As explained in Section 2.1, the interpretation is that any non-routine wage premium translates into higher effective units of non-routine labor. We treat labor as a homogeneous input, and a worker with better "non-routine" skills (such as creativity or interpersonal skills) as a worker embodying more effective units of labor. Thus, in our calibration exercise, we

²²In particular, we can compute *x* using the Cobb-Douglas assumption: $\ln x_t = \frac{\ln y_t - \alpha_t \ln k_{n,t}}{1 - \alpha_t}$, where α_t is our time varying estimate of the non-ICT capital income share.

will use as data inputs our measure of "effective" labor units derived in Section 2.1 and depicted in Figure 4.

We further allow for the possibility of labor augmenting technological progress, by measuring aggregate labor inputs as $l_t = \hat{L}_t \exp(\lambda_a t)$, where \hat{L} is the amount of effective units of labor—employment, holding worker attributes such as education, etc. fixed (see Section 2.1)—and $\lambda_a \ge 0$ is the rate of labor-augmenting technological progress. We do not allow for capital-augmenting technological progress, primarily because we would like to restrict attention to models in which the economy converges to a balanced growth path.

We therefore have seven parameters to calibrate: η , θ , γ_r , γ_{nr} , σ_r , σ_{nr} and λ_a . We impose the parametric restrictions of the CES framework, in particular that γ_i , $\eta \in [0, 1]$ and σ_i , $\theta \leq 1$. In addition, we impose that $\sigma_{nr} \leq 0.9$: this restriction is necessary because, using our calibration strategy, there is always a trivial solution in which all inputs are perfect substitutes. However, this trivial solution seems unlikely and is not supported by empirical evidence (e.g. Gaggl and Wright, 2015; Autor, 2015). In addition, we restrict attention to the case of positive labor augmenting technological progress, $\lambda_a \geq 0$.

To calibrate the above parameters, we target seven moments. Let $k_{c,i,t}$ denote the levels of the capital stock in sector i = r, nr, such that, given the levels of routine and non-routine labor in the data and a given set of parameters, the marginal product of ICT capital is equalized across sectors. Using this notation, our seven target moments are:

- 1. The change in log output net of non-ICT capital income ($\ln x$) between 1968 and 2012, given $l_{r,t}$, $l_{nr,t}$ and $k_{c,t}$
- 2. The change in the routine labor income share between 1968 and 2012, given $l_{r,t}$, $l_{nr,t}$ and $k_{c,t}$
- 3. The mean of the routine labor income share between 1968 and 2012, given $l_{r,t}$, $l_{nr,t}$ and $k_{c,t}$
- 4. The change in the ICT income share between 1968 and 2012, given $l_{r,t}$, $l_{nr,t}$ and $k_{c,t}$
- 5. The mean of the ICT income share between 1968 and 2012, given $l_{r,t}$, $l_{nr,t}$ and $k_{c,t}$
- 6. The mean of the routine employment share between 1968 and 2012, given l_t , $k_{c,t}$ and $k_{c,i,t}$
- 7. The change in the routine employment share between 1968 and 2012, given l_t , $k_{c,t}$, and $k_{c,i,t}$

Note that our targets are long run trends and average levels and we do not use any moments re-

lated to shorter-term fluctuations or co-movements in our variables. Appendix E provides further details of the calibration procedure.

The first five moment conditions are calculated taking labor allocations as given. Given routine and non-routine labor inputs, we calculate the optimal allocation of ICT capital across x_r and x_{nr} , and use it to derive equilibrium output and income shares, under the assumption that factors are paid their marginal products. The last two moment conditions are calculated using the optimality of labor allocation across routine and non-routine occupations: given simulated ICT capital allocations ($\tilde{k}_{c,i,t}$) and aggregate labor, we compute simulated labor inputs by requiring that wages are equalized.

Since we are targeting seven moment conditions to calibrate seven parameters, absent any parametric restrictions, we would expect an exact solution. Unfortunately, it turns out that the parametric restrictions are quite restrictive, so that we get a unique solution which is not exact.

To illustrate the role of parametric restrictions, consider for example the case of a single moment condition and a single unknown parameter, a. Assume that we have a model that suggests that E(z) = a, where z is some random variable of which we observe realizations $z_1, ..., z_n$. Using the sample mean as a stand-in for E(z), we can then specify our objective as: $\min_a(a - \frac{1}{n}\sum_{i=1}^n z_i)^2$. If a is unrestricted, then there is a unique exact solution, $a = \frac{1}{n}\sum_{i=1}^n z_i$, which sets the objective to 0. However, if we restrict $a \ge 2$ and we happen to observe $\frac{1}{n}\sum_{i=1}^n z_i = 1$, we get a corner solution, a = 2, which does not set the objective to 0. Note that this situation does not imply that the model is incorrect or misspecified; rather, it may reflect a situation in which the sample mean is not equal to the true mean of the distribution (for example, if n is small or if the errors $\{z_i - E(z)\}$ are systematically correlated, which is likely whenever $\{z_i\}$ is time series data).

The above example illustrates that, when parametric restrictions are binding, there could be a unique corner solution. Indeed, we find a unique corner solution in which $\gamma_r = 0$ and $\lambda_a = 0$. Figure 11 illustrates our model's fit conditional on this solution. Note that, given $\gamma_r = 0$, σ_r is irrelevant and $x_r = l_r$. We are thus left with the functional form in equation (13), similar to the specification in Krusell et al. (2000). The estimate $\lambda_a = 0$ is consistent with the view that technological improvement in the production of computing power, resulting in higher ICT capital accu-



Notes: The blue dashed lines plot the data between 1968 and 2012, and the red solid lines are calibrated equilibrium paths, given aggregate labor and aggregate ICT capital. Note that all income shares are reported as a fraction of x rather than y in this exercise and therefore are not directly comparable to those reported in Figure 1.

mulation, was the main driver of productivity growth over the past few decades (e.g., Colecchia and Schreyer, 2002; Basu et al., 2003; Bloom et al., 2012; Acemoglu and Autor, 2011).

Table 3 presents the full set of calibrated core parameters. We find an elasticity of substitution between ICT capital and non-routine labor to be around 1.2. Interestingly, this magnitude is similar to the elasticity of substitution between aggregate capital and aggregate labor estimated by Karabarbounis and Neiman (2014). Relative to aggregate capital, ICT is more substitutable with labor, but relative to labor, non-routine labor is less substitutable with capital. The magnitudes are therefore not inconsistent. We calibrate the elasticity of substitution between routine labor and x_{nr} to be around 5.7.

Table 3: Calibration of Core Parameters

Parameter	Calibratio	on	
γ_r	0		
γ_{nr}	0.0119		
σ_{nr}	0.2269	(EOS: 1.2276)	
η	0.5494		
θ	0.847	(EOS: 5.6721)	
λ_a	0	. ,	

Notes: The parameters γ_{nr} and η depend on the scaling of ICT capital. EOS is the implied elasticity of substitution.

5. Counterfactual Analysis

We embed our estimates in a neoclassical growth model to study the effects of the declining ICT price. Formally, we consider a representative agent model, in which an infinitely lived household solves the following optimization problem:

$$\max_{c_t,k_{i,t+1},l_{i,t}}\sum_{t=0}^{\infty}\beta^t u(c_t)$$

subject to $l_{r,t} + l_{n,t} = 1$, $y_t = Ak_n^{\alpha} x_t^{1-\alpha}$, equation (13), and the budget constraint

$$c_t + (1+\lambda_l) \sum_i p_{i,t} k_{i,t+1} = y_t + \sum_i p_{i,t} (1-\delta_{i,t}) k_{i,t}.$$
(14)

All variables are in per-worker terms: c_t is consumption per-worker, $k_{i,t}$ is the capital stock perworker of type i, y_t is output per-worker and $l_{r,t}$ and $l_{nr,t}$ are the employment shares in routine and non-routine occupations. The parameter λ_l captures the growth rate of labor.

We allow for two sources of exogenous variation. The first and most relevant is the decline in the price of ICT capital, $p_{c,t}$. As explained by Karabarbounis and Neiman (2014), while prices are conceptually endogenous variables, in this context the price of ICT can be interpreted as the real transformation rate of output into ICT capital. Thus, a declining ICT price captures technological progress in the production of ICT capital.²³ The second source of exogenous variation that we

²³Formally, assume that $k_{c,t+1} - (1 - \delta_c)k_{c,t} = a_t y_t^c$, where y_t^c is the output spent on ICT investment. A competitive market for the production of ICT capital goods would imply an ICT price of $p_c = \frac{1}{a}$.

Parameter	Calibration			
β	0.9747			
u(c)	$\ln(c)$			
δ_n	0.0594			
α	0.3476			
λ_l	0.028			
$p_{c,t}$	$\in [0.28, 1.62]$	(BEA data)		
$\delta_{c,t}$	$\in [0.142, 0.208]$	(BEA data)		

Table 4: Calibration of Remaining Parameters

Notes: The ICT price and depreciation rates are measured directly based on the BEA's fixed asset accounts.

consider is time variation in the depreciation rate of ICT capital, which, as we document, is quite substantial (see Figure 8).²⁴

To assess the implications of the declining ICT price, we can therefore compare our baseline simulation with a counterfactual scenario in which the ICT price remains constant at its 1968 level.²⁵ Our simulations are based on the calibration of the core model parameters (Table 3) and the values reported in Table 4. The value of β , the discount factor, is calibrated based on our estimates of the returns to capital.²⁶ Given our focus on long run trends, we assume log utility, or a unitary inter-temporal elasticity of substitution, which is consistent with the wide range of empirical estimates (see, for example, the discussion and references listed by Guvenen, 2006). Moreover, log utility is a useful benchmark for expositional purposes, as income and substitution effects cancel out.

For the depreciation rate of non-ICT capital we take the average value of our BEA-based estimates. Likewise, we calibrate the non-ICT capital intensity, α , using the average of our estimated non-ICT capital income share, ignoring any short term fluctuations around this average. As initial conditions for capital, we use the 1968 capital stocks in the US, to make our simulated transitions

²⁴We do this to improve the fit in our calibration but this does not materially affect our main results concerning the impact of the declining ICT price.

²⁵While we start our simulation in 1968 purely due to data availability reasons, it is a curious coincidence that 1968 is also the founding year of Intel Corporation, which produced the first commerically available single-chip central processing unit (CPU) in 1971.

²⁶The steady state Euler equation implies that $\beta(1 + r) = 1$, where *r* is the return to capital. We assume that the returns to capital before 1980 are roughly at their steady state level and calibrate β based on this relation.



Notes: The figures contrast the simulated effects of the declining ICT price (red line), a counterfactual simulation in which the ICT price is held constant at its 1968 level (orange line), and corresponding counter parts in the data (blue line). The shaded area indicates the "in-sample" period.

comparable with the data.

Figure 12 illustrates the simulated transition paths, and Table 5 presents a comparison of the long run implications in the baseline simulation and the counterfactual. At around 1980, the price of ICT capital starts falling rapidly. In response, agents accumulate ICT capital, and the path of ICT starts diverging from its counterfactual. Since non-ICT capital is complementary to x, the accumulation of ICT capital raises the returns to non-ICT capital, and the stock of non-ICT capital increases as well. As a result of ICT and non-ICT capital accumulation, output increases relative to its counterfactual. The higher capital stocks also raise the marginal product of labor, resulting

Variable	Change in SS relative to counterfactual				
Quantities					
Output	+12%				
Consumption	+9.4%				
non-ICT Capital	+12%				
ICT Capital	+263%				
Real Wage	+8.5%				
Income Shares					
Labor Share	-2.26%				
Routine Share	-12.26%				
Non-routine Share	+10%				
ICT Share	+10%				
Welfare Gain	+3.6%				

Table 5: Effects of the Declining ICT Price

Notes: The welfare gain is calculated as the change in consumption that leaves the representative agent indifferent with respect to the counterfactual price path at t = 0 (corresponding to the year 1968). This is different from the steady state welfare gain (which is here equal simply to the percent increase in steady state consumption), as it takes into account the transitional costs of capital accumulation and discounts steady state consumption gains.

in a higher equilibrium wage rate.

The accumulation of ICT capital leads to a divergence in the income shares of routine and nonroutine labor, roughly consistent with the magnitudes observed in the data. The net effect on the aggregate labor income share is a 2.6% decline, countered by an increase in the ICT capital income share.

In terms of welfare, the model suggests that the declining ICT price leads to a welfare gain that is equivalent to a permanent increase in consumption of 3.6%. The welfare gains are lower than the steady state consumption gains of 9.4%, for two reasons: first, the welfare figure takes into account the transitional costs associated with capital accumulation. Second, as illustrated by the transitional dynamics, there are no significant consumption gains until 1990; thus, from the perspective of 1968, consumption gains are heavily discounted. The steady state output gains of 12% reflect larger ICT and non-ICT capital stocks, together with appropriate adjustments in the allocation of labor across routine and non-routine occupations.

5.1. The Case of Homogeneous Capital and Labor Inputs

It is instructive to compare our results with an alternative calibration that treats labor and capital as homogenous inputs, rather than disaggregating labor into routine and non-routine, and capital into ICT and non-ICT. In particular, we use a CES production function with capital and labor inputs, as in Karabarbounis and Neiman (2014):

$$Y = A(\gamma K^{\sigma} + (1 - \gamma)L^{\sigma})^{\frac{1}{\sigma}}$$
(15)

We choose σ to match an elasticity of substitution between capital and labor of 1.28—the point estimate in Karabarbounis and Neiman (2014). We calibrate γ to best match the capital and labor income shares, given the levels of capital per worker in the US during our sample period, and calibrate *A* to match the levels of output given the observed levels of the capital stock.

We construct a capital aggregate using 2005 prices. Formally, we define $k_t = p_{c,2005}k_{c,t} + p_{n,2005}k_{n,t}$. The resulting price index is $p_t = \frac{1}{p_{n,t}} \frac{p_{c,t}k_{c,t}+p_{n,t}k_{n,t}}{k_t}$, using our finding that $p_{n,t}$ is roughly equal to the GDP deflator. The aggregate price index declines relative to the GDP deflator by about 25%, consistent with magnitudes discussed in Karabarbounis and Neiman (2014). Note that Karabarbounis and Neiman (2014) study the steady state implications of the declining capital price, and focus on the world as a whole. Since we focus on the US and are interested in the transitional dynamics, our results are not directly comparable (in particular, our welfare analysis takes into account the transitional costs associated with accumulating capital to reach its higher steady state level). The purpose of this comparison is to illustrate the differences and similarities in the quantitative implications resulting from our disaggregation of capital and labor inputs.

Table 6 summarizes the results. Consistent with the findings of Karabarbounis and Neiman (2014), both simulations imply that the declining investment price results in a 2-3% decline in the labor income share, about half of the observed decline in the aggregate labor income share. However, the steady state implications and the welfare implications are somewhat more modest when considering the disaggregated specification. The increases in steady state output, consumption, and wages are over 1.5 times higher when assuming homogeneous inputs, and the corresponding welfare gain is nearly 40% higher.

Variable	riable Change in SS relative to counterfactual					
Output	+20%	(+12%)				
Consumption	+17%	(+9.4%)				
Wage	+15.5%	(+8.5%)				
Labor Income Share	-2.93%	(-2.26%)				
Welfare Gain	+5%	(+3.6%)				

Table 6: Effects of the Declining ICT Price (Homogeneous Labor & Capital)

Notes: The welfare gain is defined as in Table 5. Numbers in parentheses are the corresponding values in Table 5

The higher welfare gains are not obvious a-priori: on the one hand, since the ICT price was declining faster than the price of other capital (Figure 8), a homogeneous capital input is associated with a milder price decline and hence lower implied welfare gains. On the other hand, since ICT represents a relatively small share of capital expenditure, a specification with homogeneous capital inputs will overstate the importance of declining capital prices, implying higher welfare gains. Our results suggest that the latter bias dominates quantitatively.

6. Concluding Remarks

The discussion of the social costs and benefits of automation can roughly be divided into two issues: the effects on aggregate consumption and the distributional implications. The general consensus is that, while there is likely a positive effect on aggregate consumption, there are some adverse distributional implications. The net welfare implications depend on the relative quantitative importance of the two.

In this paper we set out to derive an upper bound for the quantitative benefits of automation, by employing a representative agent model in which redistribution of income is frictionless. We assess the welfare gains from automation at about 3.6% of consumption per worker, from the perspective of 1968. This welfare gain takes into account the transitional costs of accumulating ICT capital, but not the R&D costs associated with generating the decline in the ICT capital price, which we treat as exogenous. Therefore, 3.6% is, in some sense, an upper bound—though, it should be noted that if instead of 1968 we use 1980 as a base year, the welfare gains would be

higher since future consumption gains are discounted at a lower rate.

Even though the representative agent framework is unable to account for distributional costs, it is informative regarding potential distributional implications. Our analysis suggests that while ICT may have a large effect on the distribution of labor income, it has only a moderate effect on the distribution of income between capital and labor.

Our results further suggest that automation is unlikely to be the sole cause of the declining labor income share. In particular, our measurement suggests that only half of the decline in the labor income share is directly countered by an increase in the ICT capital income share. The remainder is due to an increase in the non-ICT capital income share, particularly in the post-2001 period and primarily driven by housing. This suggests that, in order to fully understand the sources of the declining labor income share, it is necessary to study the mechanisms driving changes in non-ICT capital income, in particular residential capital income. We leave this challenge for future work.

References

Acemoglu D. 1999. Changes in unemployment and wage inequality: An alternative theory and some evidence. *American Economic Review* 89: 1259–1278.

URL http://ideas.repec.org/a/aea/aecrev/v89y1999i5p1259-1278.html

- Acemoglu D, Autor D. 2011. Skills, Tasks and Technologies: Implications for Employment and Earnings, volume 4 of Handbook of Labor Economics, chapter 12. Elsevier, 1043–1171. URL http://ideas.repec.org/h/eee/labchp/5-12.html
- Acemoglu D, Autor D, Dorn D, Hanson GH, Price B. 2014. Return of the Solow Paradox? IT, Productivity, and Employment in US Manufacturing. American Economic Review 104: 394–99. URL http://ideas.repec.org/a/aea/aecrev/v104y2014i5p394–99.html
- Akerman A, Gaarder I, Mogstad M. 2013. The Skill Complementarity of Broadband Internet. IZA Discussion Papers 7762, Institute for the Study of Labor (IZA). URL http://ideas.repec.org/p/iza/izadps/dp7762.html
- Atkinson AB, Piketty T, Saez E. 2011. Top Incomes in the Long Run of History. Journal of Economic Literature 49: 3–71.
- Autor D, Dorn D, Hanson GH. 2013. Untangling Trade and Technology: Evidence from Local Labor Markets. IZA Discussion Papers 7329, Institute for the Study of Labor (IZA).

URL http://ideas.repec.org/p/iza/izadps/dp7329.html

URL http://ideas.repec.org/a/aea/jeclit/v49y2011i1p3-71.html

Autor DH. 2015. Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives* **29**: 3–30.

URL http://www.aeaweb.org/articles.php?doi=10.1257/jep.29.3.3

Autor DH, Dorn D. 2013. The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review* **103**: 1553–97.

URL http://ideas.repec.org/a/aea/aecrev/v103y2013i5p1553-97.html

Autor DH, Katz LF, Kearney MS. 2008. Trends in u.s. wage inequality: Revising the revisionists. *The Review of Economics and Statistics* **90**: 300–323.

URL http://ideas.repec.org/a/tpr/restat/v90y2008i2p300-323.html

Autor DH, Levy F, Murnane RJ. 2003. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal* of *Economics* **118**: 1279–1333.

URL http://ideas.repec.org/a/tpr/qjecon/v118y2003i4p1279-1333.html

Basu S, Fernald JG, Oulton N, Srinivasan S. 2003. The Case of the Missing Productivity Growth: Or, Does Information Technology Explain why Productivity Accelerated in the US but not the UK? NBER Working Papers 10010, National Bureau of Economic Research, Inc.

URL http://ideas.repec.org/p/nbr/nberwo/10010.html

Bloom N, Sadun R, Van Reenen J. 2012. Americans Do IT Better: US Multinationals and the Productivity Miracle. *American Economic Review* **102**: 167–201.

URL http://ideas.repec.org/a/aea/aecrev/v102y2012i1p167-201.html

Bridgman B. 2014. Is Labor's Loss Capital's Gain? Gross versus Net Labor Shares. BEA Working Papers 0114, Bureau of Economic Analysis.

URL http://ideas.repec.org/p/bea/wpaper/0114.html

- Brynjolfsson E, Hitt LM. 2003. Computing productivity: Firm-level evidence. *The Review of Economics and Statistics* **85**: 793–808. URL http://ideas.repec.org/a/tpr/restat/v85y2003i4p793–808.html
- Caselli F, Feyrer J. 2007. The marginal product of capital. *The Quarterly Journal of Economics* **122**: 535–568. URL http://ideas.repec.org/a/tpr/qjecon/v122y2007i2p535-568.html
- Champagne J, Kurmann A. 2012. Reconciling the divergence in aggregate us wage series. Working Paper. URL http://www.andrekurmann.com/files/wp_files/CESpaper_27July2012.pdf

Christensen LR, Jorgenson DW. 1969. The measurement of us real capital input, 1929–1967. Review of Income and Wealth 15: 293–320.

- Colecchia A, Schreyer P. 2002. ICT Investment and Economic Growth in the 1990s: Is the United States a Unique Case? A Comparative Study of Nine OECD Countries. *Review of Economic Dynamics* 5: 408–442. URL http://ideas.repec.org/a/red/issued/v5y2002i2p408-442.html
- Cortes GM, Jaimovich N, Nekarda CJ, Siu HE. 2014. The micro and macro of disappearing routine jobs: A flows approach. Working Paper 20307, National Bureau of Economic Research. URL http://www.nber.org/papers/w20307

Dorn D. 2009. Essays on inequality, spatial interaction, and the demand for skills. Dissertation 3613, University of St. Gallen.

Eden M, Gaggl P. 2015. Do poor countries really need more IT ? the role of relative prices and industrial composition. Policy Research Working Paper Series 7352, The World Bank.

URL http://ideas.repec.org/p/wbk/wbrwps/7352.html

Elsby MW, Hobijn B, Sahin A. 2013. The decline of the us labor share. Brookings Papers on Economic Activity .

- Gaggl P, Wright GC. 2015. A short run view of what computers do: Evidence from a uk tax incentive. Working paper, UNC Charlotte. URL http://belkcollegeofbusiness.uncc.edu/pgaggl/wp-content/uploads/sites/36/2014/09/UK_ICT_ 05-08-2015_WP.pdf
- Goldin C, Katz L. 2008. *The Race Between Education and Technology*. Harvard University Press. ISBN 9780674028678. URL https://books.google.com/books?id=mcYsvvNEUYwC
- Goldin C, Katz LF. 1998. The Origins Of Technology-Skill Complementarity. *The Quarterly Journal of Economics* **113**: 693–732. URL http://ideas.repec.org/a/tpr/qjecon/v113y1998i3p693-732.html
- Goos M, Manning A. 2007. Lousy and lovely jobs: The rising polarization of work in britain. *The Review of Economics and Statistics* **89**: 118–133.

URL http://ideas.repec.org/a/tpr/restat/v89y2007i1p118-133.html

- Goos M, Manning A, Salomons A. 2009. Job polarization in europe. *American Economic Review* **99**: 58–63. URL http://ideas.repec.org/a/aea/aecrev/v99y2009i2p58-63.html
- Guvenen F. 2006. Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective. Journal of Monetary Economics 53: 1451–1472. URL http://ideas.repec.org/a/eee/moneco/v53y2006i7p1451–1472.html
- Hall RE, Jorgenson DW. 1967. Tax policy and investment behavior. *The American Economic Review* 57: pp. 391–414. ISSN 00028282. URL http://www.jstor.org/stable/1812110
- Hennessy J, Patterson D, Asanović K. 2012. Computer Architecture: A Quantitative Approach. Computer Architecture: A Quantitative Approach. Morgan Kaufmann/Elsevier. ISBN 9780123838728. URL https://books.google.com/books?id=v3-lhVwHnHwC
- Jones CI, Kim J. 2014. A schumpeterian model of top income inequality. Working Paper 20637, National Bureau of Economic Research. URL http://www.nber.org/papers/w20637
- Jorgenson DW. 1995. Productivity: Postwar US economic growth, volume 1. Mit Press.
- Jorgenson DW, Vu K. 2007. Information Technology and the World Growth Resurgence. *German Economic Review* 8: 125–145. URL http://ideas.repec.org/a/bla/germec/v8y2007ip125-145.html
- Karabarbounis L, Neiman B. 2014. The global decline of the labor share. *The Quarterly Journal of Economics* **129**: 61–103. URL http://gje.oxfordjournals.org/content/129/1/61.abstract

- Krusell P, Ohanian LE, Ríos-Rull JV, Violante GL. 2000. Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis. *Econometrica* 68: 1029–1054. URL http://ideas.repec.org/a/ecm/emetrp/v68y2000i5p1029-1054.html
- Levy F, Murnane RJ. 1992. U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations. *Journal of Economic Literature* **30**: 1333–81.

URL http://ideas.repec.org/a/aea/jeclit/v30y1992i3p1333-81.html

- O'Mahony M, Van Ark B. 2003. *EU productivity and competitiveness: an industry perspective: can Europe resume the catching-up process?* Office for Official Publications of the European Communities Luxembourg.
- Piketty T, Saez E. 2003. Income inequality in the united states, 1913-1998. The Quarterly Journal of Economics 118: 1–39. URL http://ideas.repec.org/a/tpr/qjecon/v118y2003i1p1-39.html
- Rognlie M. 2015. Deciphering the fall and rise in the net capital share. In *Brookings Papers on Economic Activity*. Conference Draft. URL http://www.brookings.edu/~/media/projects/bpea/spring-2015/2015a_rognlie.pdf
- Ruggles S, Alexander JT, Genadek K, Goeken R, Schroeder MB, Sobek M. 2010. Integrated public use microdata series: Version 5.0 [machine-readable database]. Technical report, University of Minnesota, Minneapolis. URL https://cps.ipums.org/cps/
- vom Lehn C. 2015. Labor Market Polarization, the Decline of Routine Work, and Technological Change: A Quantitative Evaluation. 2015 Meeting Papers 151, Society for Economic Dynamics. URL http://ideas.repec.org/p/red/sed015/151.html

Appendix A. Construction of Capital Income Shares

This appendix outlines our approach to measure income shares directly from the BEA's nominal current cost values of detailed asset categories. As discussed in Section 3, our measurement exercise relies on two assumptions: First, no-arbitrage in capital markets requires the *gross return* (8) to be equalized across all types of capital, which can be written more compactly as

$$(1 + \pi_{i,t})\frac{R_{i,t}}{P_{i,t}} + CG_{i,t} = R_t$$
(A.1)

where $(1 + \pi_{i,t}) \equiv P_{i,t+1}/P_{i,t}$ is gross price inflation and $CG_{i,t} = (1 + \pi_{i,t})(1 - \delta_t)$ are capital gains for capital of type *i*.

Second, we impose constant returns to scale in aggregate production (equation (7)), and thus

the following condition must hold:

$$s_{K,t}P_tY_t = \sum_{i \in I} R_{i,t}K_{i,t}$$
(A.2)

where $s_{K,t}$ is a measure of the income share paid to reproducible capital. Equations (A.1) and (A.2) directly imply the following expression for the relative price adjusted rental rates for each capital type $i = \{c, n\}$:

$$PR_{c,t} = \frac{R_{c,t}}{P_{c,t}} = \frac{P_{n,t}K_{n,t}}{\tilde{P}_t\tilde{K}_t}(CG_{n,t} - CG_{c,t}) + (1 + \pi_{n,t})s_{K,t}\frac{P_tY_t}{\tilde{P}_t\tilde{K}_t}$$
(A.3)

$$PR_{n,t} = \frac{R_{n,t}}{P_{n,t}} = \frac{P_{c,t}K_{c,t}}{\tilde{P}_{t}\tilde{K}_{t}}(CG_{c,t} - CG_{n,t}) + (1 + \pi_{c,t})s_{K,t}\frac{P_{t}Y_{t}}{\tilde{P}_{t}\tilde{K}_{t}}$$
(A.4)

where $\tilde{P}_t \tilde{K}_t \equiv (1 + \pi_{n,t}) P_{c,t} K_{c,t} + (1 + \pi_{c,t}) P_{n,t} K_{n,t}$. Based on (A.3) and (A.4) we then compute the income shares for each type of capital as

$$s_{c,t} = PR_{c,t} \frac{P_{c,t}K_{c,t}}{P_t Y_t}$$
(A.5)

$$s_{n,t} = PR_{n,t} \frac{P_{n,t}K_{n,t}}{P_t Y_t}$$
(A.6)

where we use the following data inputs:

- 1. $(1 + \pi_{i,t})$ ICT/non-ICT prices based on the BEA's fixed asset accounts
- 2. $(1 \delta_{i,t})$ depreciation rates from the BEA's fixed asset accounts accounts
- 3. $P_{i,t}K_{i,t}$ nominal capital stocks from the BEA's fixed asset accounts accounts
- 4. $P_t Y_t$ nominal GDP from the BEA
- 5. $s_{K,t}$ measured as the reciprocal of the official BLS labor share

Appendix B. CPS Outoing Rotation Groups vs. March Supplements

Figure B.13 illustrates a comparison of the disaggregate labor income shares based on two alternative measures of earnings. The details for the data construction are discussed in Section 2. The



Figure B.13: Labor's Income Share: MORG vs. MARCH

Notes: Occupation specific income shares are based on CPS earnings data from the annual march supplement (1968 and after) and the monthly outgoing rotation groups (MORG, starting in 1979) extracts and rescaled to match the aggregate income share in the Non-Farm Business Sector (BLS). The underlying earnings data for both series are top-code adjusted using Piketty and Saez's (2003) updated estimates of the income distribution (PS). The MORG series shows annual averages of monthly data that was seasonally adjusted using the U.S. Cenusus X11 method. The graphs are constructed analogously to panel A in Figure 1. For details see Section 2.

thick lines are the shares reported in panel A of Figure 1 and are based on the march supplements in the CPS (MARCH). The thin lines use usual weekly earnings from the outgoing rotation groups (MORG).

Appendix C. Labor Expenditure Embodied in Net Exports

Panels A and B of Figure C.14 illustrate the numerator of equation (6) for goods and services separately. Panel A makes clear that the US is predominantly trading sectors that are very capital intensive and therefore the labor expenditure embodied in goods trade is very small (compared to GDP). Panel Breveals that the US is a net exporter of labor embodied in services, predominantly in goods that are very non-routine labor intensive. Panel C aggregates goods and services, and panel D plots $s_{NX,\ell,t}$ for both routine and non-routine labor.



Figure C.14: Labor Expenditure Embodied in Net Exports

Notes: Panels A and B illustrate the labor expenditures embodied in net exports of goods and services, respectively. Panels C and D illustrate the sum of A and B both in millions of USD and as a percent of GDP.

Appendix D. Price-Quantity Decomposition: Housing



Notes: Panel A graphs implicit price deflators by capital type relative to the GDP deflator, which were constructed directly from the BEA's detailed fixed-asst accounts. The BEA GDP deflator is taken from FRED. Panel B depicts asset-specific depreciation rates constructed directly from the BEA's fixed asset accounts. The dashed vertical lines indicate the year 1968.

Appendix E. Calibrating the aggregate production function

We specify our objective as follows:

$$\max_{\eta,\theta,\sigma_r,\sigma_{nr},\gamma_r,\gamma_{nr},\lambda_a} \sum_{i=1}^7 w_i (\tilde{g}_i - g_i)^2$$
(E.1)

Where g_i are moment conditions and w_i are their corresponding weights. We use tildes to denote simulated variables and moments. Using p_t^x to denote the price of the input x, the moment conditions are as follows:

$$g_1 = \ln(x_{2012}) - \ln(x_{1968}) \tag{E.2}$$

$$g_2 = \frac{1}{2012 - 1968 + 1} \sum_{t=1968}^{2012} \frac{w_{r,t} L_{r,t}}{p_t^x x_t}$$
(E.3)

$$g_3 = \frac{1}{2012 - 1968 + 1} \sum_{t=1968}^{2012} \frac{R_{c,t} K_{c,t}}{p_t^x x_t}$$
(E.4)

$$g_4 = \frac{1}{2012 - 1968 + 1} \sum_{t=1968}^{2012} \frac{L_{r,t}}{L_t}$$
(E.5)

$$g_5 = \frac{w_{r,2012}L_{r,2012}}{p_{2012}^x x_{2012}} - \frac{w_{r,1968}L_{r,1968}}{p_{1968}^x x_{1968}}$$
(E.6)

$$g_6 = \frac{R_{c,2012}K_{c,2012}}{p_{2012}^x x_{2012}} - \frac{R_{c,1968}K_{c,1968}}{p_{1968}^x x_{1968}}$$
(E.7)

$$g_7 = \frac{L_{r,2012}}{L_{2012}} - \frac{L_{r,1968}}{L_{1968}} \tag{E.8}$$

The corresponding weights are chosen as 1 for all moment conditions, except for the ones relating to ICT capital shares, which receive weights of 40. These weights are chosen so that the calibrated production function generates roughly the same quality of fit for all moment conditions. The higher weights assigned to the ICT capital conditions are necessary because the ICT capital share is substantially smaller than the other shares considered in the other moment conditions.

The target values of the moment conditions are computed as follows. The value of x_t is computed as $\ln(x_t) = \frac{\ln(Y_t) - \alpha_{k,t} \ln(K_{n,t})}{1 - \alpha_{k,t}}$, where $K_{n,t}$ is the non-ICT capital stock and $\alpha_{k,t}$ is the non-ICT capital income share. The income shares in g_2 , g_3 , g_5 and g_6 are computed as output shares divided by $1 - \alpha_{k,t}$ (since, for example, $\frac{w_r L_r}{p^{x_x}} = \frac{w_r L_r}{y} \frac{y}{p^{x_x}}$, and, since the expenditure on input x is a fraction $1 - \alpha_k$ of output, this amounts to $\frac{s_r}{1 - \alpha_k}$).

The target employment shares in conditions g_4 and g_7 are calculated as the share of routine labor income out of total labor income. As explained in the text, this has the interpretation of the equalization of wages per effective unit of labor across routine and non-routine occupations.

To calculate the simulated moments, we use the following procedure. The inputs for the simulation include the set of parameters, as well as data on aggregate labor supply; aggregate ICT supply; and the supply of routine and non-routine labor. The procedure begins by calculating the distribution of ICT capital between the production of x_r and x_{nr} , given labor allocations of $L_{r,t}$ and $L_{nr,t}$, and given an aggregate ICT capital stock of $K_{c,t}$. This results in simulated values $\tilde{K}_{c,r,t}$ and $\tilde{K}_{c,nr,t}$. The value of \tilde{x}_t and the simulated income shares are calculated using the $\tilde{K}_{c,r,t}$ and $\tilde{K}_{c,nr,t}$ as ICT capital inputs, and $L_{r,t}$ and $L_{nr,t}$ as labor inputs, under the assumption that factors are paid their marginal products. Note that the computations here take routine and non-routine labor inputs from the data. This procedure results in the simulated moments $\tilde{g}_1 - \tilde{g}_3$ and $\tilde{g}_5 - \tilde{g}_6$.

The next step calculates the equilibrium allocation of labor, given an aggregate labor supply of L_t , and ICT capital stocks of $\tilde{K}_{c,r,t}$ and $\tilde{K}_{c,nr,t}$. This yields simulated values for routine and non-routine labor, $\tilde{L}_{r,t}$ and $\tilde{L}_{n,t}$. These simulated values are then used for calculating the simulated moments \tilde{g}_7 and \tilde{g}_4 .

Note that though the information contained in routine and non-routine labor income shares is the same as the information contained in routine and non-routine employment shares, the procedure utilizes different equilibrium conditions to target income and employment shares. In particular, the moments relating to income shares are computed taking routine and non-routine employment as given, while the moments relating to employment shares require the optimality of the allocation of labor across routine and non-routine occupations.

Since the model is highly non-linear, our objective function has many local minima. To obtain a global minimum, we use a Newton-type convergence algorithm, starting from over 10,000 randomly drawn initial conditions.

Appendix F. Classification of ICT and Non-ICT Assets

	Share of	f Aggregate Ca	pital (%)	Average Growth in Share (%)		
ICT Assets	1960-1980	1980-2000	2000-2013	1960-1980	1980-2000	2000-2013
EP20: Communications	2.73	3.91	3.39	2.87	1.78	-2.63
ENS3: Own account software	0.24	0.75	1.56	27.26	6.68	2.58
ENS2: Custom software	0.11	0.61	1.40	34.82	8.49	2.06
EP34: Nonelectro medical instruments	0.35	0.76	1.08	4.87	2.97	2.30
EP36: Nonmedical instruments	0.51	0.92	0.92	0.62	2.41	-1.08
ENS1: Prepackaged software	0.02	0.33	0.83	32.28	14.63	-1.04
EP35: Electro medical instruments	0.11	0.36	0.66	7.25	3.43	4.28
EP1B: PCs	0.00	0.31	0.45		12.12	0.96
RD23: Semiconductor and other component manufacturing	0.05	0.23	0.43	6.58	8.21	2.75
RD22: Communications equipment manufacturing	0.26	0.21	0.27	3.27	0.89	0.24
EP31: Photocopy and related equipment	0.53	0.75	0.26	6.75	-2.11	-7.70
EP1A: Mainframes	0.19	0.36	0.24	24.00	1.91	-4.97
EP1H: System integrators	0.00	0.03	0.23		42.85	3.45
RD24: Navigational and other instruments manufacturing	0.05	0.19	0.22	3.20	5.78	-1.59
EP1D: Printers	0.07	0.22	0.19	20.75	7.20	-9.76
EP1E: Terminals	0.02	0.14	0.16	71.14	5.48	-4.62
EP1G: Storage devices	0.00	0.17	0.12		7.55	-9.55
EP12: Office and accounting equipment	0.48	0.32	0.12	-3.09	-5.00	-6.13
RD40: Software publishers	0.00	0.05	0.09		16.91	-1.13
RD21: Computers and peripheral equipment manufacturing	0.16	0.09	0.07	3.68	-3.07	-0.60
RD25: Other computer and electronic manufacturing, n.e.c.	0.01	0.01	0.02	0.91	3.24	-0.34
EP1C: DASDs	0.09	0.13	0.00	30.38	-36.26	-78.36
EP1F: Tape drives	0.06	0.03	0.00	22.77	-40.33	-186.06

Table F.7: ICT Assets

Notes: The data are drawn from the BEA's detailed fixed asset accounts. ICT assets are defined as BEA asset codes starting with EP, EN, RD2, or RD4. Notice that the EP category incorporates two assets that are strictly speaking likely not ICT: EP34 and EP36. Our results do not critically hinge on these two assets and we therefore stick with the more standard BEA aggregation of EP. Assets are ranked by their average share in aggregate capital during 2000-2013. The share of aggregate capital is the value of each individual asset, as estimated by the BEA at current cost, as a fraction of the value of all assets in Tables F.7 and F.8. Panel A reports averages of these shares for three time periods. Panel B reports the average annual growth in these shares over same three time periods.

Table F.8: Non-ICT Assets

B Average Growth in Share (%)

	A. Average Share of Aggregate Capital (%)		ate Capital (%)	B. Average Growth in Share (%)		
Non-ICT Assets	1960-1980	1980-2000	2000-2013	1960-1980	1980-2000	2000-2013
SOO1: Office	4.75	6.99	8.28	1.59	2.14	0.12
SI00: Manufacturing	8.27	8.07	6.96	0.22	-0.44	-1.39
SM01: Petroleum and natural gas	3.91	3.63	5.20	0.12	-2.30	6.54
SU30: Electric	6.44	5.48	4.73	-0.18	-1.79	0.89
SC03: Multimerchandise shopping	2.21	2.77	3.05	1.46	0.69	0.59
EI50: General industrial equipment	3.42	3.40	2.95	0.32	-0.60	-0.75
SB31: Hospitals	1.70	2.52	2.86	3.69	1.06	0.51
SU20: Communication	2.80	2.51	2.66	0.70	-1.09	1.79
SC02: Other commercial	1.55	2.00	2.47	1.43	1.41	0.38
SB41: Lodging	1.58	1.83	2.32	1.40	1.76	0.47
EI60: Electric transmission and distribution	2.87	2.50	2.16	-0.79	-0.64	-0.21
EI40: Special industrial machinery	2.63	2.52	1.95	-0.36	-0.07	-3.33
SB20: Educational and vocational	1.56	1.36	1.92	-0.26	0.68	2.78
SC01: Warehouses	1.18	1.40	1.89	0.54	1.57	1.03
ET30: Aircraft	1.17	1.59	1.79	5.36	0.94	0.37
EO80: Other	1.09	1.44	1.75	2.30	1.18	0.71
EO12: Other furniture	1.37	1.67	1.74	-0.45	1.68	-1.35
SB42: Amusement and recreation	1.84	1.67	1.70	-0.46	0.54	-1.44
SN00: Farm	3.48	2.46	1.69	-0.63	-2.66	-2.11
RD11: Pharmaceutical and medicine manufacturing	0.23	0.65	1.68	4.70	6.63	5.22
SU40: Gas	2.70	1.89	1.66	-1.66	-1.72	0.35
SC04: Food and beverage establishments	1.19	1.51	1.58	1.47	0.84	-0.51
SB10: Religious	2.11	1.57	1.51	-0.70	-0.73	-0.82
EI30: Metalworking machinery	2.32	2.10	1.49	0.46	-1.25	-3.31
EI11: Light trucks (including utility vehicles)	0.93	1.05	1.42	0.45	2.48	-2.56
RDOM: Other manufacturing	1.47	1.50	1.19	0.00	0.04	-1.13
E120: Autos	1.80	1.67	1.14	-1.79	-0.33	-4.27
E112: Other trucks, buses and truck trailers	1.59	1.46	1.02	0.54	-1.43	-3.04
SU11: Other railroad	4.14	1.86	0.97	-4.62	-3.94	-3.77
SU12: Track replacement	2.66	1.39	0.97	-4.22	-2.69	-1.34
SOO2: Medical buildings	0.57	0.83	0.95	1.55	1.90	0.25
EO40: Other construction machinery	1.16	1.02	0.94	1.79	-2.11	1.37
AETU: Theatrical movies	0.97	0.67	0.86	-4.28	2.37	0.00
AE20: Long-lived television programs	0.69	0.79	0.80	0.39	1.91	-0.55
EU60. Service industry machinery	1.14	0.94	0.85	-1.01	-0.49	0.00
EITZ. Other labricated metals	1.42	1.20	0.75	0.83	-3.88	-0.19
RD60: All other nonmanulacturing, n.e.c.	0.09	0.75	0.72	3.04	0.20	-3.34
SD32. Special calle	0.41	1.02	0.71	3.70	1.59	-0.74
EO30: Other agricultural machinery	1.58	1.09	0.00	-1.57	-4.52	-0.73
El21: Stoom ongines	0.75	0.57	0.05	0.01	2.01	-0.74
SU50: Petroleum ninelines	0.75	0.57	0.40	-2.76	-3.01	1 69
AE30: Books	0.35	0.37	0.43	-0.37	1.02	-0.54
ET40: Shins and hoats	0.40	0.68	0.40	-0.91	-3.72	-1 18
BD92: Other nonprofit institutions	0.02	0.00	0.38	3 94	2 25	0.23
RD31: Motor vehicles and parts manufacturing	0.22	0.33	0.36	0.45	0.79	-4 90
SM02: Mining	0.31	0.40	0.35	2 12	-1.30	1 93
BD12: Chemical manufacturing ex pharma and med	0.59	0.10	0.34	0.04	-0.98	-1.58
EO50: Mining and oilfield machinery	0.54	0.39	0.31	0.73	-5 74	7 21
SO01: Water supply	0.23	0.28	0.28	0.22	1 25	-0.60
EO21: Farm tractors	0.69	0.44	0.28	-0.08	-4.65	-0.28
SO02: Sewage and waste disposal	0.24	0.29	0.28	0.26	1.16	-1.21
RD32: Aerospace products and parts manufacturing	0.33	0.41	0.25	2.09	-0.28	-2.54
AE40: Music	0.21	0.20	0.20	0.14	1.37	-4.21
RD70: Scientific research and development services	0.00	0.06	0.19		11.29	4.20
SO04: Highway and conservation and development	0.15	0.18	0.18	0.27	1.20	-0.44
SB43: Air transportation	0.15	0.15	0.17	0.74	1.03	-0.69
AE50: Other entertainment originals	0.18	0.17	0.17	-1.71	1.53	-2.58
SU60: Wind and solar	0.00	0.01	0.16		12.44	25.22
EO72: Other electrical	0.12	0.19	0.14	2.74	-0.20	-1.74
SO03: Public safety	0.14	0.10	0.11	-0.96	0.17	-0.55
RD60: Computer systems design and related services	0.00	0.03	0.11		22.31	2.68
EO11: Household furniture	0.16	0.13	0.10	0.19	-2.30	-1.38
EO22: Construction tractors	0.28	0.18	0.09	0.39	-5.15	-3.98
EI22: Internal combustion engines	0.09	0.08	0.08	-0.80	-1.19	0.43
SB44: Local transit structures	0.51	0.17	0.07	-6.06	-5.15	-4.84
EI11: Nuclear fuel	0.03	0.10	0.06	27.52	-3.93	-1.86
RD91: Private universities and colleges	0.03	0.04	0.06	0.95	2.71	3.68
SOMO: Mobile structures	0.05	0.07	0.05	0.35	1.29	-3.77
SB46: Other land transportation	0.04	0.03	0.05	-0.72	1.04	2.80
RU50: Financial and real estate services	0.00	0.02	0.05	0.00	22.33	-0.88
EU/I: Housenoid appliances	0.09	0.05	0.03	-2.09	-4.08	-2.35
SD45. Uther transportation	0.02	0.02	0.02	-0.70	0.56	-0./1

Notes: The data are drawn from the BEA's detailed fixed asset accounts. Assets are ranked by their average share in aggregate capital during 2000-2013. The share of aggregate capital is the value of each individual asset, as estimated by the BEA at current cost, as a fraction of the value of all assets in Tables F.7 and F.8. Panel A reports averages of these shares for three time periods. Panel B reports the average annual growth in these shares over same three time periods.