

Technological Change and Earnings Inequality in the U.S.: Implications for Optimal Taxation*

Pedro Brinca[†] João B. Duarte[‡] Hans A. Holter[§] João G. Oliveira[¶]

May 9, 2022

[Download latest version](#)

Abstract

Since 1980 there has been a steady increase in earnings inequality alongside rapid technological growth in the U.S. economy. To what extent does technological change explain the observed increase in earnings dispersion? How does it affect the optimal progressivity of the tax system? To answer these questions, we develop an incomplete markets model with occupational choice, in which occupations are heterogeneous along two dimensions: cognitive complexity and routine task intensity. We find that technological transformation can fully account for the increase in earnings dispersion between 1980 and 2015. The main driver is the rising relative wage of non-routine cognitive occupations, which benefit the most from complementarity with capital. In isolation, increasing earnings inequality strengthens the case for redistributive policies. However, with occupational choice, we find that a significant drop in the tax system's progressivity is socially optimal because it leads workers into higher-paid occupations, reducing the wage dispersion in equilibrium.

Keywords: Income Inequality, Taxation, Technological Change, Automation

JEL Classification: E21; H21; J31.

*We thank Árpád Ábrahám, Daron Acemoglu, David Autor, Vasco Botelho, Juan Dolado, Loukas Karabarbounis, Nick Kozeniauskas, Musa Orak, Lee Ohanian, Cezar Santos, Pedro Teles, Gianluca Violante, seminar participants at the Lisbon Macro Group, the EUI Economics Department and participants at the PEJ 2018, the 2019 Winter Meeting of the North-American Econometric Society, and the 2019 Lubra-Macro conference for their helpful comments and suggestions. Pedro Brinca and João Duarte are grateful for financial support from the Portuguese Science and Technology Foundation (FCT), grants number SFRH/BPD/99758/2014, UID/ECO/00124/2013 and UID/ECO/00145/2013, POR Lisboa (LISBOA-01-0145-FEDER-007722 and Social Sciences Data Lab, Project 22209), and POR Norte (Social Sciences Data Lab, Project 22209). João Duarte is also thankful for his CEECIND/03227/2018 FCT grant. João Oliveira is grateful for financial support from the Portuguese Science and Technology foundation, grant number SFRH/BD/138631/2018.

[†]Center for Economics and Finance at Universidade of Porto and Nova School of Business and Economics.

[‡]Nova School of Business and Economics.

[§]University of Oslo and Nova School of Business and Economics.

[¶]Nova School of Business and Economics. Corresponding author. E-mail address: joaooliveira.2011@novasbe.pt.

1 Introduction

Earnings inequality in the U.S. has increased steadily since 1980, see Figure 1 (left panel)¹. Alongside the increase in inequality there has also been technological progress. The right panel of Figure 1 displays a rapid fall in the relative price of equipment investment goods, which can be viewed as reflecting investment-specific technological change such as cheaper access to computing power and storage (Krusell et al., 2000; Karabarbounis and Neiman, 2014). In this paper, we answer the following questions: (i) to what extent does investment-specific technological change explain the observed increase in earnings inequality? (ii) how does it affect the optimal progressivity of the tax and transfer system?

To answer these questions, we develop a life-cycle, incomplete markets, overlapping generations model featuring uninsurable idiosyncratic earnings risk, technological change, a detailed tax system, and occupational choice. Central to our approach is that we adopt the framework of Autor et al. (2003) where occupations differ in terms of the nature of the tasks that are being performed. There are four main categories of tasks: Non-routine cognitive (NRC), non-routine manual (NRM), routine cognitive (RC) and routine manual (RM). Households choose an occupation at the beginning of their work lives based on an idiosyncratic cost of acquiring the necessary skills and on the distribution of future earnings in each profession.

We find that occupation-biased technological change can account fully for the increase in earnings inequality between 1980 and 2015. The main driver is the rising relative wage of non-routine cognitive occupations, which benefit the most from complementarity with capital. Thus, investment-specific technological change stands as a major engine behind the growth of earnings dispersion, together with occupation-biased technological change.


We find the optimal value of our measure of tax progressivity in 1980 to be 0.15, whereas in 2015 a value of 0.05 is optimal.² The main mechanisms driving this result are the shifts in occupational choice resulting from a drop in progressivity, which raise the wages in lower-paid occupations; the increase in the return rates of capital relative to 1980, which weaken the insurance role of a progressive tax system;³ and the accumulation of more productive capital, which raises output per capita.⁴

¹In fact Figure 1 shows a steady increase since 1970. However, our paper will for the most part focus on the period from 1980 to 2015 due to availability of other data.

²Indeed there has been some reduction in tax progressivity in the U.S: since 1980, see Wu (2021)

³See Jordà et al. (2019) for evidence of higher return rates on wealth in the U.S.

⁴We use a non-linear functional form as in Benabou (2002) Heathcote et al. (2019), $y_a = 1 - \theta_0 y^{-\theta_1}$ where θ_0 and θ_1 define the level and progressivity, respectively.



Note: The pre-tax earnings Gini is computed from the CPS for employed workers. Description of the data is provided on section 3. The relative price of investment is computed as the ratio between equipment investment prices from the BEA and the BLS urban consumer price index.

Figure 1: Inequality and ISTC.

Although there is more earnings inequality in 2015, which creates additional incentives for redistribution, the positive effects that high-earning occupations have on the wages of low-earning occupations are stronger and dominate redistributive motives. This is, in part, because occupational choice is endogenous in our model, which weakens the impact of rising income dispersion in the considerations of the social planner.

We estimate an aggregate production function, where each of the four occupations, capital equipment and capital structures are inputs. The extent to which labor demand and wages in each of these occupation categories will affect the wage distribution is determined by their respective roles in the production function, by latent occupation-biased technological change, and, in particular, by their complementarity with capital equipment. In simple terms: The fall in the price of equipment investment goods spurs capital accumulation and creates increased demand for workers in occupations with tasks which are more complementary with capital relative to those that are less so. If there are barriers to mobility between occupations and different costs of entry, the rise in labor demand for some occupations creates a wage premium relative to workers in other occupations. In our model, barriers to mobility are modeled as an idiosyncratic cost of acquiring the necessary skills at labor market entry, with no additional mobility allowed during a household's work life.

The rest of the paper is organized as follows. Section 2 contains a brief survey of the related literature. In Section 3, we discuss the stylized facts that underlie our modeling choices. In Section 4, we describe the model. In Section 5 we estimate our aggregate

production function. Section 6 is devoted to calibrating our model. In Section 7, we present our quantitative results on inequality and optimal taxation. Section 8 concludes.

2 Related literature

This paper is related to the literature investigating the impact of technological change on inequality. Our first contribution is to expand on the seminal paper by [Krusell et al. \(2000\)](#), who document the role of skill-biased technological change and capital-skill complementarity in the skill premium, by specifying and estimating an aggregate production function with labor inputs based on occupations rather than the levels of education of the workforce. To do this, we use the occupation taxonomy of [Autor et al. \(2003\)](#), who study the effect of computerization on changes in employment by occupation categories. They posit that some occupations have a prevalence of tasks which involve complex problem-solving and interaction (so-called non-routine tasks) which are very costly or nigh impossible to automate using a pre-defined set of instructions. We use the cross-walk developed by [Cortes et al. \(2020\)](#) to map tasks into occupation codes, in order to calculate equilibrium quantities of labor input by occupation category. [Eden and Gaggl \(2018\)](#) also estimate an aggregate production function for the U.S. using the routine/non-routine paradigm and investigate the welfare implications of investment-specific technological change for the welfare of a representative agent. In contrast, our work uses the four task dimensions postulated by [Autor et al. \(2003\)](#) and also allows for labor-augmenting technological change at the occupation level, which are important for our conclusion that workers at the bottom of the distribution have enjoyed wage growth relative to the center of the wage distribution as a result of technological change. While our model draws heavily on the task-based framework used in the literature ([Acemoglu and Autor, 2011](#); [Acemoglu and Restrepo, 2018](#), [Moll et al., 2019](#)) we do not model tasks explicitly. We thus forego a more detailed characterization of the production process in favor of the ability to measure inputs more accurately, enabling the estimation a more abstract version of the production technology.

This paper is also related to the literature on optimal taxation. Our contribution is to assess the impact of technological change on optimal progressivity in an incomplete markets model with occupational choice. This is similar to the focus of the recent work by [Guerreiro et al. \(2021\)](#), who study capital taxation in a model with the possibility of automation of tasks, and endogenous skill/occupation choice. Our contribution is distinct from theirs in that we broaden the analysis to include the cognitive/manual dimensions of tasks, and focus on the progressivity of the labor income tax schedule. Like

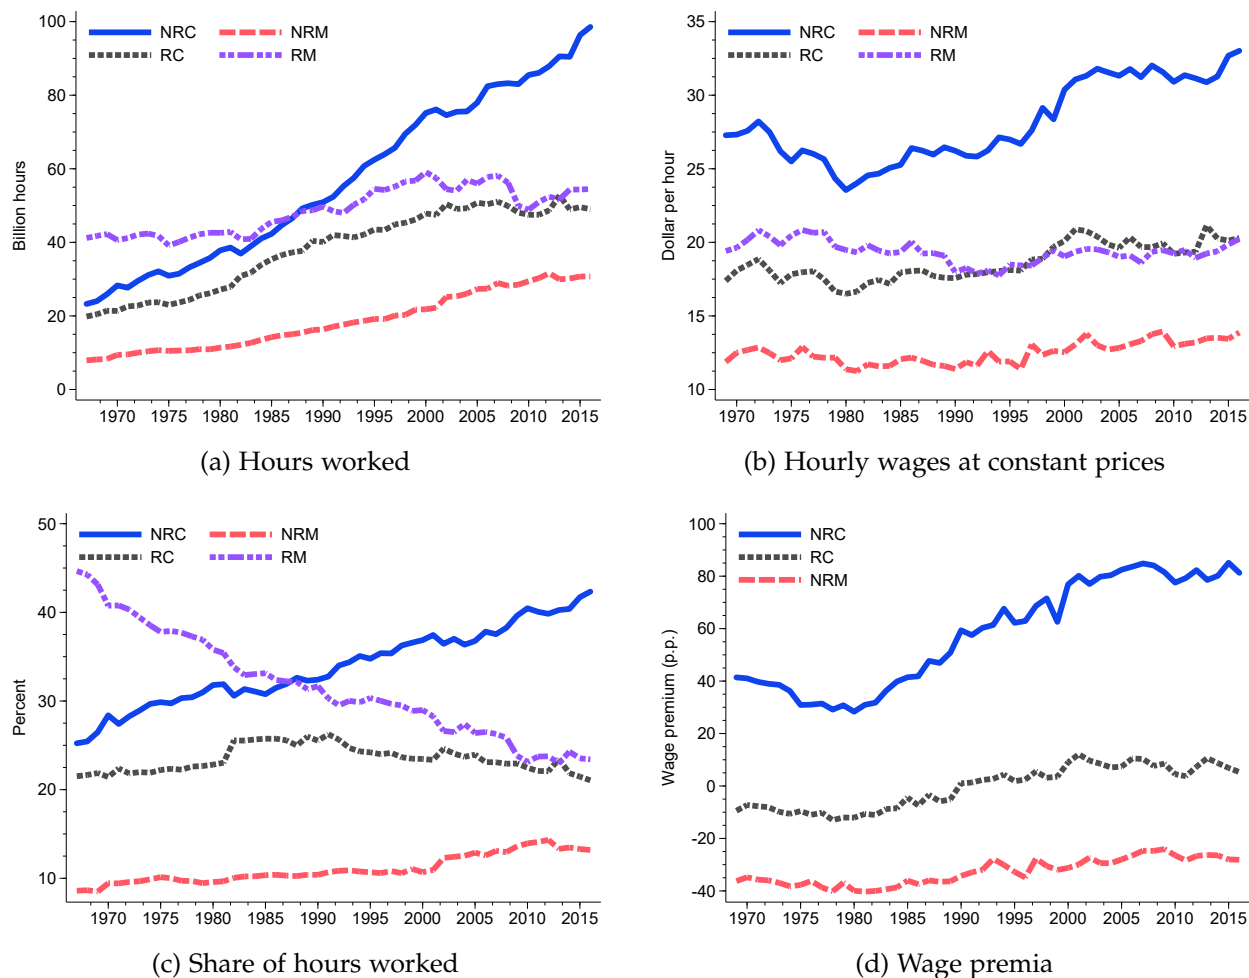
them, however, we assume that older generations cannot change occupations, which is in line with the evidence provided by [Cortes et al. \(2020\)](#), who argue that the fall in routine employment in the U.S. has been primarily caused by declining inflow rates among younger workers. The closest paper to our own is [Heathcote et al. \(2020\)](#), who study the impact of technological change on optimal progressivity in an incomplete markets model with occupational choice. However, their focus is on college education and skill-biased technological change, while our paper takes an occupation-based approach.

To dilute the perceived social cost of these trends several policies have been suggested, including a proposal famously put forth by Bill Gates to tax robots and “even slow down the speed [of automation]” ([Delaney, 2017](#)). The issue of optimal capital taxation has been discussed extensively since the seminal papers by [Chamley \(1986\)](#) and [Lucas \(1990\)](#), who find that the optimal rate of capital taxation is zero in the steady state. In contrast, [Aiyagari \(1995\)](#), using an incomplete markets model with borrowing constraints, determines that the optimal income tax on capital income is positive, even in the long run.

3 Motivating Facts

Data. Our analysis of the U.S. labor market is carried out along the two dimensions proposed by [Autor et al. \(2003\)](#) to classify occupations: (i) whether the main tasks are more susceptible to automation (routine) or less (non-routine); and (ii) the nature of the tasks involved, i.e., whether they are predominantly cognitive or manual. This classification system yields four mutually exclusive occupation groups: non-routine cognitive (NRC), non-routine manual (NRM), routine cognitive (RC) workers and routine manual (RM). We use data from the Census Bureau Current Population Survey (CPS), spanning the period from 1968 to 2016, to study how quantities and prices have changed since the late 1960s for each of these groups.

We use the Annual Social and Economic Supplement (ASEC) from the March CPS survey available from [Flood et al. \(2018\)](#), which contains data on yearly earnings and hours worked in the previous calendar year. The CPS employs the US Census Bureau 2010 occupation classification system, and we use the cross-walk table of [Cortes et al. \(2020\)](#) to categorize each worker into one of the aforementioned classes. This cross-walk is based on the so-called “consensus” classification scheme of [Acemoglu and Autor \(2011\)](#). The population of interest is the set of non-military, non-institutionalized individuals aged 16 to 70, excluding the self-employed and farm sector workers. See Appendix [A](#) for additional details on data treatment. These data are used to construct



Note: Wage premia are obtained as the log difference between the constant composition average wage of each occupation category. Groups for wages are constructed with using a constant composition of individual observable characteristics (experience, education, etc). Data source is the CPS Annual Social and Economic Supplement. See sections A and B of the Appendix for details.

Figure 2: Employment and wages by occupation category.

time series on employment and wages by occupation category. To calculate wage premia we use the method of [Krusell et al. \(2000\)](#), as described in Appendix B.

Facts. Figure 2 shows the evolution of employment and wages for the selected occupation categories. From 1968 to 2016, hours worked increased roughly five-fold in the NRC category, three-fold in NRM, doubled in RC, and nearly stagnated in RM.

There are three main takeaways: (i) the strong performance of NRC workers compared to other groups and, in particular, to RM workers; (ii) the growth of cognitive worker groups relative to manual; (iii) the rise of non-routine cognitive wage premium.

The central hypothesis in this paper is that one of the main drivers of the increase

in inequality since the 1980s has been the discriminating effect that investment specific technological change has had on these four groups due to its diverse interaction with each labor variety. This reasoning is similar to that of [Krusell et al. \(2000\)](#), [Karabarbounis and Neiman \(2014\)](#), [Acemoglu and Restrepo \(2017\)](#), and [Eden and Gaggli \(2018\)](#).

This choice was made due to the quantitative importance of ISTC for the long-run growth of output per hours worked in the U.S. economy, originally estimated to be 60% in [Greenwood et al. \(1997\)](#), as well as its potential to disrupt labor market conditions. Indeed, [Krusell et al. \(2000\)](#) used a model of capital-skill complementarity and ISTC to study the increased wage dispersion in the U.S economy and are able to track the progress of the skill premium. Similarly to [Acemoglu and Restrepo \(2017\)](#) and [Eden and Gaggli \(2018\)](#), we view the process of ISTC as akin to increased automation of routine tasks in the economy. However, we focus on the wage premium rather than on worker displacement in this paper.

Central to investment-specific technological change are the falling prices of capital goods, which can be interpreted as evidence of increasing productivity in the investment goods sector. As an illustration of this interpretation, consider that in the 1950s a computer was leased for 200,000 per month in inflation-adjusted 2010 dollars, plus the costs of the staff and energy required to operate it.⁵ Today, any computer or smartphone equipped with microprocessors costs a fraction of that price and is able to deliver a processing speed which is many million times that of a large-scale computer in the 1950s.⁶ To get a sense of the scale of technological change, the CPU of a Play Station 2 is 1,500 times faster than the guidance computer on Apollo 11, while the Apple iPhone4 is 4,000 times faster.

Is there reason to believe that this source of growth has a uniform impact across labor markets? [Krusell et al. \(2000\)](#) argue that this is not the case. Using aggregate U.S. data they estimate the parameters for a CES production function where capital, skilled and unskilled labor are embedded. They find that capital is a gross complement with skilled labor and a gross substitute for unskilled labor. Therefore, secular growth is skill-biased and is able to reproduce the rise in the skill premium observed in the U.S. since the start of the 1980s, highlighting the importance of worker training for productivity and inequality. Both [Karabarbounis and Neiman \(2014\)](#) and [Eden and Gaggli \(2018\)](#) depart from similar hypotheses in building their frameworks.

⁵Source: http://ethw.org/Early_Popular_Computers,_1950_-_1970.

⁶Not to mention holding a much larger quantity of information: in 1956, IBM's 305 RAMAC disk could hold 5 MB of information, while the computer on which this paper was written has a total of 4.78 TB in hard drive memory.

4 A model of labor market inequality and technological change

Our model is of the Bewley-Aiyagari-Hugget variety:⁷ An incomplete markets economy with overlapping generations of heterogeneous agents and partially uninsurable idiosyncratic risk that generates both an income and a wealth distribution. Households derive utility from non-durable consumption and leisure.

Prior to entering the labor market, households choose their occupation type based on an idiosyncratic cost of acquiring the necessary skills to perform it. For tractability, we assume that this decision is irreversible and mutually exclusive, and determines from which labor market the household will draw its wage over the course of its lifetime.⁸ After labor market entry, households face an idiosyncratic stream of earnings in the form of wages, and make joint decisions about consumption, savings and hours worked.

For the production side of the economy, we draw heavily on the modeling strategy of [Krusell et al. \(2000\)](#) and [Karabarbounis and Neiman \(2014\)](#). There are three final goods sectors in the economy: Consumption goods, structure capital goods, and an equipment capital goods sector. This formulation allows us to express the price of equipment goods as a function of the level of technology in that sector relative to the consumption goods sector, which is the formulation that [Krusell et al. \(2000\)](#) adopt in order to incorporate investment-specific technological change.

The centerpiece of the model is the production function of the intermediate input, which uses a combination from the different occupation and capital types in order to produce final goods. We build on the production function introduced by [Greenwood et al. \(1997\)](#) and extend it in order to encompass a total of four labor varieties: Non-routine cognitive, non-routine manual, routine cognitive, and routine manual.

Technological progress, in the form of total factor productivity growth, occupation-biased technological change, and investment-specific technological change, affects capital and labor demand and, thereby, occupation wage premia. This framework creates a rich interaction between capital accumulation, technological change, and the wages of different occupations and allows us to map the dynamics in these variables into earnings inequality measures.

One key mechanism driving wage dispersion in this economy is investment-specific

⁷See [Bewley \(2000\)](#), [Aiyagari \(1994\)](#), and [Hugget \(1993\)](#).

⁸[Cortes et al. \(2020\)](#) show evidence that the main driver of decline in routine employment has been a reduction in inflow rates rather than increases in outflows. This is consistent with our assumption of inability of changing occupation type in the middle of working life, in spite of changing wage premia in other occupation types.

technological change: As equipment prices fall, firms substitute away from routine manual labor to equipment capital and other types of labor which are more complementary with capital. Shifting demand of labor varieties coupled with limited labor mobility produces changes in wage premia over time.

In this section, we describe the household problem, the production side of the economy, and the equilibrium concept.

4.1 Demographics

We assume the economy is populated by a set of $J - 1$ overlapping generations, as in [Brinca et al. \(2016\)](#). A period in the model to correspond to one year. Thus, j , the household's age, varies between 0 (for age 20 households) and 80 (for age 100 households).

Prior to joining the labor market, agents must make an irreversible and mutually exclusive occupation choice, deciding which labor market will determine their wages over the course of their lives. Thus, a household i draws idiosyncratic utility, κ_{io} , from acquiring the necessary skills to join occupation type $o \in O = \{\text{NRC}, \text{NRM}, \text{RC}, \text{RM}\}$. This term can be viewed as the personal cost (or benefit, if positive) of the process of acquiring skills to perform the tasks associated with a given occupation type, such as the effort (or joy) from studying in the case of cognitive occupations, for example.

We assume that κ_{io} follows a type 1 extreme value distribution, H_o , with location parameter $\mu_{\kappa,o}$ and scale parameter $\sigma_{\kappa,o}$ in the tradition of discrete choice modeling of [McFadden \(1973\)](#).⁹ Households choose the occupation where total utility is highest:

$$\tilde{V}_{io} = \kappa_{io} + V_o, \tag{1}$$

where V_o is the expected discounted lifetime utility from choosing occupation type o , κ_{io} is the idiosyncratic utility draw for occupation o . Assuming $\sigma_{\kappa,o} = 1, \forall o \in O$, this formulation allows us to write the probability of choosing an occupation o before κ_{io} is known as:

$$p_o = \frac{e^{\mu_o + V_o}}{\sum_{l \in O} e^{\mu_l + V_l}}. \tag{2}$$

As a result, equation 2 is also the closed form expression for the employment share of occupation o .¹⁰ Other than occupation, households differ in the value of their persis-

⁹Concretely, this formulation is the same as that used for unordered multinomial models where discrete choices are modeled as outcomes from an additive random utility model. See [Cameron and Trivedi \(2005\)](#) for a detailed exposition.

¹⁰In order to find V_o for each occupation, we calibrate and solve a version of the model where occupations are randomly assigned in such a way that we match the employment weights of each occupation

tent idiosyncratic productivity shock, u_{ij} , permanent ability, a_i , and asset holdings, h_{ij} . Working age agents have to choose how much to work, n_{ij} , how much to consume, c_{ij} , and how much to save, h_{ij+1} , to maximize utility.

After retiring at age 65, households face an age-dependent probability of dying, $\pi(j)$, dying with certainty at age 100. $s_j = 1 - \pi_j$ defines the age-dependent probability of surviving, and so, at any given period, using a law of large numbers, the mass of retired agents of age $j \geq 45$ is equal to $S_j = \prod_{t=45}^{t=j} s_t$.

A fraction of households leave unintended bequests which are redistributed in a lump-sum manner between the households that are currently alive, denoted by Γ . We include a bequest motive in this framework to make sure that the age distribution of wealth is empirically plausible, as in [Brinca et al. \(2019\)](#).

Retired households make consumption and saving decisions and receive a retirement benefit, $\Psi(a_i)$. For simplicity, we assume that the public retirement benefit is constant until the agent's death and is equal to a fraction, ψ_{ss} , of the average earnings of an agent with permanent ability a_i at age $j = 44$ working $1/3$ of its time. ψ_{ss} is such that the Social Security system breaks even in equilibrium.

4.2 Preferences

The period utility function, $u(c_{ij}, n_{ij})$, depends on consumption, c_{ij} , and labor supply, $n_{it} \in (0, 1]$, and is given by:¹¹

$$u(c_{it}, n_{it}) = \log c_{ij} - \chi \frac{n_{ij}^{1+\eta}}{1+\eta}, \quad (3)$$

where η is the inverse Frisch elasticity of labor supply. Log utility over consumption ensures that the existence of a balanced-growth path. The utility function of retired households has one extra term, as they gain utility from the bequest they leave to living generations:

$$D(b_{ij+1}) = \varphi \log(b_{ij+1}). \quad (4)$$

type in 1980. Employment shares used are computed from CPS data and are: $p_{\text{NRC}} = 0.302$, $p_{\text{NRM}} = 0.109$, $p_{\text{RC}} = 0.243$. We then compute the expected utility for each occupation type, V_o , at age 20 which we use to solve and calibrate the version of the model with occupational choice.

¹¹We assume that labor disutility depends only on the level of supply, not on occupation type.

where b_{ij+1} is the level of liquid savings of household i . The expected discounted lifetime utility of household i after occupational choice is given by:

$$V = \mathbb{E}_0 \left[\sum_{j=1}^J \beta^{j-1} [S_j u(c_{ij}, n_{ij}) + (1 - S_j) D(b_{ij+1})] \right], \quad (5)$$

where β is the discount factor and $S_j = 1$ for $j < 45$.

4.3 Labor income

Labor productivity depends on three elements which determine the amount of efficiency units each household is endowed with in each period: Age, j , permanent ability, a_i , and the idiosyncratic productivity shock, u_{ij} , which we assume follows an AR(1) process:

$$u_{ij} = \rho_u u_{ij-1} + \varepsilon_{ij}, \quad \varepsilon_{ij} \sim N(0, \sigma_\varepsilon^2). \quad (6)$$

Thus, household i 's wage at age j is given by:

$$w_i(j, o, a_i, u_{ij}) = w_o e^{\gamma_0 + \gamma_1 j + \gamma_2 j^2 + \gamma_3 j^3 + a_i + u_{ij}}, \quad (7)$$

where γ_1 , γ_2 and γ_3 are estimated directly from the data to capture the age profile of wages, and γ_0 is set such that the age polynomial is equal to zero at age 20 in the model. Households' labor income also depends on the wage per efficiency unit of labor w_o , $o \in O \equiv \{\text{NRC}, \text{NRM}, \text{RC}, \text{RM}\}$, where o is the labor variety supplied by the household and chosen at the beginning of the work life. Permanent ability is assigned at labor market entry and has variance $\sigma_{a,o}$ which depends on the occupation, in order to match within-occupation earnings dispersion. Appendix D describes implementation of this procedure in the numerical algorithm.

4.4 Technology

In this framework, three competitive final goods sectors exist: Consumption goods, structure investment goods, and equipment investment goods. Each is produced is produced by transforming a single intermediate input using a linear production technology. All payments are made in the consumption good, which is the numeraire.

The consumption good is obtained by transforming a quantity $Z_{c,t}$ of intermediate input into output, which is then sold at price $p_{c,t}$ to both households and the government.

The transformation technology is:

$$C_t + G_t = Z_{c,t}, \quad (8)$$

where $Z_{c,t}$ is the quantity of input, purchased at $p_{z,t}$ from a representative intermediate goods firm. Given that the consumption good is competitively produced, its price equals the marginal cost of production:

$$p_{c,t} = 1 = p_{z,t}. \quad (9)$$

Likewise, structure investment good firms produce output with a similar technology:

$$X_{s,t} = Z_{s,t}, \quad (10)$$

and therefore $p_{s,t} = 1$. The production of $X_{e,t}$, the equipment investment good, uses the transformation technology:

$$X_{e,t} = \frac{Z_{e,t}}{\zeta_t}, \quad (11)$$

where $Z_{e,t}$ is the quantity of input z used in the production of the final equipment good. $1/\zeta_t$ is the level of technology used in the production of $X_{e,t}$ relative to the final consumption good. As ζ_t declines, the relative productivity in assembling the equipment good increases. We assume that ζ_t evolves exogenously. We obtain the price of the equipment good from the zero profit condition:

$$p_{e,t} = \zeta_t p_{z,t} = \zeta_t, \quad (12)$$

where $\zeta_t = p_{e,t}/p_{c,t}$ is interpreted as the relative price of the equipment good.

A representative intermediate goods firm produces $Z_{c,t} + Z_{s,t} + Z_{e,t}$ using a constant returns to scale technology in capital and labor inputs, $y_t = F(K_{s,t}, K_{e,t}, N_{NRC,t}, N_{NRM,t}, N_{RC,t}, N_{RM,t})$, where $K_{s,t}$ is structure capital and $K_{e,t}$ is capital equipment. The firm rents structure capital at rate $r_{s,t}$, equipment at r_t^e and each labor variety at $w_{o,t}$, $o \in O$. Aggregate demand measured in terms of the consumption good, $Y_t = C_t + G_t + X_{s,t} + \zeta_t X_{e,t}$, factor prices, and the price of the intermediate good $p_{z,t}$ are taken as given. The firm chooses capital and labor inputs each period in order to maximize profits:

$$\Pi_{z,t} = p_{z,t} y_t - r_{s,t} K_{s,t} - r_{e,t} K_{e,t} - \sum_{o \in O} w_{o,t} N_{o,t}, \quad (13)$$

subject to:

$$y_t = Z_{c,t} + Z_{s,t} + Z_{e,t} = C_t + G_t + X_{s,t} + \xi_t X_{e,t} = Y_t. \quad (14)$$

This setup implies that $Z_{c,t} = C_t + G_t$, $Z_{s,t} = X_{s,t}$, $Z_{e,t} = \xi_t X_{e,t}$, and $F(\cdot) = Y_t = C_t + G_t + X_{s,t} + \xi_t X_{e,t}$. We assume that the production function of intermediate goods is Cobb-Douglas over non-equipment capital and CES over the remaining inputs:¹²

$$F(\cdot) = A_t G(\cdot) = A_t K_{s,t}^\alpha \left[\sum_{i=1}^3 \varphi_i Z_{i,t}^{\frac{\sigma-1}{\sigma}} + \left(1 - \sum_{i=1}^3 \varphi_i \right) N_{RM,t}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma(1-\alpha)}{\sigma-1}}, \quad (15)$$

$$Z_{1,t} = \left[\phi_1 K_{e,t}^{\frac{\rho_1-1}{\rho_1}} + (1 - \phi_1) N_{NRC,t}^{\frac{\rho_1-1}{\rho_1}} \right]^{\frac{\rho_1}{\rho_1-1}}, \quad Z_{2,t} = \left[\phi_2 K_{e,t}^{\frac{\rho_2-1}{\rho_2}} + (1 - \phi_2) N_{NRM,t}^{\frac{\rho_2-1}{\rho_2}} \right]^{\frac{\rho_2}{\rho_2-1}},$$

$$Z_{3,t} = \left[\phi_3 K_{e,t}^{\frac{\rho_3-1}{\rho_3}} + (1 - \phi_3) N_{RC,t}^{\frac{\rho_3-1}{\rho_3}} \right]^{\frac{\rho_3}{\rho_3-1}},$$

where A_t is total factor productivity, ϕ_i and φ_i are distribution parameters where $l = 1, 2, 3$, correspond to the occupation types NRC, NRM, and RC, respectively. ρ_l is the elasticity of substitution between capital and the nested labor variety i , and σ is the elasticity of substitution between each composite $Z_{l,t}$ and routine manual labor. Complementarity between the two inputs in $Z_{l,t}$ requires that $\rho_l < \sigma$, as in [Krusell et al. \(2000\)](#).

Each variety of labor input is measured in efficiency units, $N_{o,t} \equiv h_{o,t} \varrho_{o,t}$, where $h_{o,t}$ is the quantity of hours worked in the aggregate and $\varrho_{o,t}$ is an efficiency index representing the latent quality per hour worked in occupation type o in period t . $\varrho_{o,t}$ can be interpreted as an occupation-specific technology level, due to research and development, or as human capital accumulation.

¹²[Krusell et al. \(2000\)](#), [Karabarbounis and Neiman \(2014\)](#), and [Eden and Gaggl \(2018\)](#) use CES production functions where capital equipment is nested with all labor varieties except for unskilled/routine manual labor, which is introduced in isolation. The reason for this setup is the set of symmetry restrictions on substitution elasticities imposed by the CES functional form, as explained in [Krusell et al. \(2000\)](#). In a nutshell, this nesting form allows for complementarity between capital and differentiated labor (NRC, NRM, RC) while permitting the elasticities of substitution between routine manual labor and other labor varieties to be different. Our version is an extension of this framework with a finer breakdown over labor varieties. In estimating the production function, we use the Simulated pseudo-Maximum Likelihood Estimation (SPMLE) method proposed by [Ohanian et al. \(1997\)](#) which was also applied in [Krusell et al. \(2000\)](#). Our application is described in the next section.

Firm maximization implies that marginal products equal factor prices:¹³

$$w_{\text{NRC},t} = \Xi_t \varphi_1 \left[\phi_1 \left(\frac{K_{e,t}}{N_{\text{NRC},t}} \right)^{\frac{\rho_1-1}{\rho_1}} + (1 - \phi_1) \right]^{\frac{\sigma-\rho_1}{(\rho_1-1)\sigma}} [1 - \phi_1] Q_{\text{NRC},t}, \quad (16)$$

$$w_{\text{NRM},t} = \Xi_t \varphi_2 \left[\phi_2 \left(\frac{K_{e,t}}{N_{\text{NRC},t}} \right)^{\frac{\rho_2-1}{\rho_2}} + (1 - \phi_2) \left(\frac{N_{\text{NRM},t}}{N_{\text{NRC},t}} \right)^{\frac{\rho_2-1}{\rho_2}} \right]^{\frac{\sigma-\rho_2}{(\rho_2-1)\sigma}} [1 - \phi_2] \left(\frac{N_{\text{NRM},t}}{N_{\text{NRC},t}} \right)^{-\frac{1}{\rho_2}} Q_{\text{NRM},t}, \quad (17)$$

$$w_{\text{RC},t} = \Xi_t \varphi_3 \left[\phi_3 \left(\frac{K_{s,t}}{N_{\text{NRC},t}} \right)^{\frac{\rho_3-1}{\rho_3}} + (1 - \phi_3) \left(\frac{N_{\text{RC},t}}{N_{\text{NRC},t}} \right)^{\frac{\rho_3-1}{\rho_3}} \right]^{\frac{\sigma-\rho_3}{(\rho_3-1)\sigma}} [1 - \phi_3] \left(\frac{N_{\text{RC},t}}{N_{\text{NRC},t}} \right)^{-\frac{1}{\rho_3}} Q_{\text{RC},t}, \quad (18)$$

$$w_{\text{RM},t} = \Xi_t (1 - \varphi_1 - \varphi_2 - \varphi_3) \left(\frac{N_{\text{RM},t}}{N_{\text{NRC},t}} \right)^{-\frac{1}{\sigma}} Q_{\text{RM},t}, \quad (19)$$

$$r_{s,t} = A_t \alpha \left[\frac{K_{e,t}}{N_{\text{NRC},t}} \right]^{\alpha-1} \Lambda_t^{\frac{\sigma(1-\alpha)}{\sigma-1}}, \quad (20)$$

$$r_{e,t} = \Xi_t \left[\varphi_1 \left(\phi_1 \left[\frac{K_{e,t}}{N_{\text{NRC},t}} \right]^{\frac{\rho_1-1}{\rho_1}} + [1 - \phi_1] \right)^{\frac{\sigma-\rho_1}{(\rho_1-1)\sigma}} \phi_1 \left(\frac{K_{e,t}}{N_{\text{NRC},t}} \right)^{-\frac{1}{\rho_1}} + \varphi_2 \left(\phi_2 \left[\frac{K_{e,t}}{N_{\text{NRC},t}} \right]^{\frac{\rho_2-1}{\rho_2}} + [1 - \phi_2] \left[\frac{N_{\text{NRM},t}}{N_{\text{NRC},t}} \right]^{\frac{\rho_2-1}{\rho_2}} \right)^{\frac{\sigma-\rho_2}{(\rho_2-1)\sigma}} \phi_2 \left(\frac{K_{e,t}}{N_{\text{NRC},t}} \right)^{-\frac{1}{\rho_2}} + \varphi_3 \left(\phi_3 \left[\frac{K_{e,t}}{N_{\text{NRC},t}} \right]^{\frac{\rho_3-1}{\rho_3}} + [1 - \phi_3] \left[\frac{N_{\text{RC},t}}{N_{\text{NRC},t}} \right]^{\frac{\rho_3-1}{\rho_3}} \right)^{\frac{\sigma-\rho_3}{(\rho_3-1)\sigma}} \phi_3 \left(\frac{K_{e,t}}{N_{\text{NRC},t}} \right)^{-\frac{1}{\rho_3}} \right], \quad (21)$$

¹³Marginal products are expressed as functions of the ratios between each factor and the non-routine cognitive labor for the purpose of constructing the solution algorithm.

where¹⁴

$$\Xi_t = A_t \left[\frac{K_{s,t}}{N_{\text{NRC},t}} \right]^\alpha [1 - \alpha] \Lambda_t^{\frac{1-\sigma\alpha}{\sigma-1}}.$$

Capital laws of motion are given by:

$$K_{s,t+1} = (1 - \delta_s)K_{s,t} + X_{s,t}, \quad (22)$$

$$K_{e,t+1} = (1 - \delta_e)K_{e,t} + X_{e,t}, \quad (23)$$

where δ_s and δ_e are the depreciation rates of structures and equipment, respectively.

4.5 Government

The social security system is managed by the government and runs a balanced budget. Revenues are collected from taxes on employees and on the representative firm at rates τ_{ss} and $\tilde{\tau}_{ss}$, respectively, and are used to pay retirement benefits, Ψ .

The government taxes consumption, τ_c , and capital income, τ_k , at flat rates. The labor income tax follows a non-linear functional form as in [Benabou \(2002\)](#) and [Heathcote et al. \(2019\)](#):

$$y_a = 1 - \theta_0 y^{-\theta_1}, \quad (24)$$

where θ_0 and θ_1 define the level and progressivity of the tax schedule, respectively. y is the pre-tax labor income and y_a is after-tax labor income.¹⁵

Tax revenues from consumption, labor, and capital income taxes are used to finance public consumption, G_t , which clears the budget constraint. Denoting social security revenues by R_t^{ss} and the other tax revenues as T_t , the government budget constraint is defined as:

¹⁴Variable Λ_t is defined as:

$$\begin{aligned} \Lambda_t = & \varphi_1 \left(\phi_1 \left[\frac{K_{e,t}}{N_{\text{NRC},t}} \right]^{\frac{\rho_1-1}{\rho_1}} + [1 - \phi_1] \right)^{\frac{\rho_1(\sigma-1)}{(\rho_1-1)\sigma}} + \varphi_2 \left(\phi_2 \left[\frac{K_{e,t}}{N_{\text{NRC},t}} \right]^{\frac{\rho_2-1}{\rho_2}} + [1 - \phi_2] \left[\frac{N_{\text{NRM},t}}{N_{\text{NRC},t}} \right]^{\frac{\rho_2-1}{\rho_2}} \right)^{\frac{\rho_2(\sigma-1)}{(\rho_2-1)\sigma}} \\ & + \varphi_3 \left(\phi_3 \left[\frac{K_{e,t}}{N_{\text{NRC},t}} \right]^{\frac{\rho_3-1}{\rho_3}} + [1 - \phi_3] \left[\frac{N_{\text{RC},t}}{N_{\text{NRC},t}} \right]^{\frac{\rho_3-1}{\rho_3}} \right)^{\frac{\rho_3(\sigma-1)}{(\rho_3-1)\sigma}} + (1 - \varphi_1 - \varphi_2 - \varphi_3) \left(\frac{N_{\text{RM},t}}{N_{\text{NRC},t}} \right)^{\frac{\sigma-1}{\sigma}}. \end{aligned}$$

¹⁵See the [Holter et al. \(2014\)](#) for a detailed discussion of the properties of this tax function.

$$T_t = G_t, \quad (25)$$

$$\Psi_t \left(\sum_{j \geq 45} \Omega_j \right) = R_t^{ss}. \quad (26)$$

4.6 Asset structure

Households hold two asset types: Structures capital, $k_{s,ij}$, and equipment capital, $k_{e,ij}$. There is no investment-specific technological change in the steady state, i.e., $\tilde{\zeta}_{t+1} = \tilde{\zeta}_t = \zeta$, so we drop the time index on return rates for this exposition. Thus, the return rates must satisfy:

$$\frac{1}{\zeta} [\zeta + (r_e - \zeta \delta_e)(1 - \tau_k)] = 1 + (r_s - \delta_s)(1 - \tau_k), \quad (27)$$

which follows from non-arbitrage: Investing in equipment capital must yield the same after-tax return as investing the same amount in structures. Total assets for the consumer are defined as:

$$b_{ij} \equiv \zeta k_{e,ij} + k_{s,ij}, \quad (28)$$

4.7 Household problem

On any given period a household is defined by age, j , occupation o_i , asset position b_{ij} , permanent ability a_i , and a persistent idiosyncratic productivity shock u_{ij} . A working-age household chooses consumption, c_{ij} , work hours, n_{ij} , and future asset holdings, b_{ij+1} , to solve its problem. The household budget constraint is given by:

$$c_{ij}(1 + \tau_c) + \zeta k_{e,ij+1} + k_{s,ij+1} = [\zeta + (r_e - \zeta \delta_e)(1 - \tau_k)] k_{e,ij} \\ [1 + (r_s - \delta_s)(1 - \tau_k)] k_{s,ij} + q\Gamma + Y^N, \quad (29)$$

where Y^N is the household's labor income after social security and labor income taxes, and $q = 1/(1 + r_s(1 - \tau_k))$. Using 27, in equilibrium we can rewrite the budget constraint as:

$$c_{ij}(1 + \tau_c) + h_{ij+1} = (h_{ij} + \Gamma)[1 + r(1 - \tau_k)] + Y^N. \quad (30)$$

The household problem can be formulated recursively as:

$$V(j, h_{ij}, o_i, a_i, u_{ijt}) = \max_{c_{ij}, n_{ij}, b_{ij+1}} \left[u(c_{ij}, n_{ij}) + \beta \mathbb{E}_{u_{j+1}} [V(j+1, b_{it+1}, o_i, a_i, u_{ij+1})] \right]$$

s.t.:

$$c_{ij}(1 + \tau_c) + b_{ij+1} = (b_{ij} + \Gamma)[1 + r(1 - \tau_k)] + Y^N$$

$$Y^N = \frac{n_{ij}\omega(j, o_i, a_i, u_{ij})}{1 + \tilde{\tau}_{ss}} \left(1 - \tau_{ss} - \tau_l \left[\frac{n_{ij}\omega(j, o_i, a_i, u_{ij})}{1 + \tilde{\tau}_{ss}} \right] \right)$$

$$n_{ij} \in (0, 1], \quad b_{ij} \geq 0, \quad b_{i0} = 0 \quad \forall i, \quad c_{ij} > 0.$$

The problem of a retired household differs in three features: A positive age-dependent probability of dying, $\pi(j)$, a bequest motive $D(b_{ij+1})$, and labor income, which is replaced by constant retirement benefit depending on permanent ability, $\Phi(a_i)$. Therefore, the retired household's problem is defined as:

$$V(j, b_{ij}, a_i) = \max_{c_{ij}, b_{ij+1}} \left[u(c_{ij}, b_{ij+1}) + \beta(1 - \pi(j))V(j+1, b_{ij+1}, a_i) + \pi(j)D(b_{ij+1}) \right]$$

s.t.:

$$c_{ij}(1 + \tau_c) + b_{ij+1} = (b_{ij} + \Gamma)[1 + r(1 - \tau_k)] + \Psi(a_i)$$

$$b_{ij+1} \geq 0, \quad c_{ij} > 0.$$

4.8 Stationary recursive competitive equilibrium

$\Phi(j, b, o, a, u)$ is the measure of agents with corresponding characteristics (j, b, o, a, u) . The stationary recursive competitive equilibrium is defined by:¹⁶

1. Taking factor prices and initial conditions as given, the value function $V(j, b, o, a, u)$ and the policy functions, $o(\kappa_o)$, $c(j, b, o, a, u)$, $b'(j, b, o, a, u)$, and $n(j, b, o, a, u)$, solve the household's optimization problem.

2. Markets clear:

$$\xi K_e + K_s = \int b + \Gamma d\Phi,$$

$$N_{RM} = \varrho_{RM} \int n_{RM} d\Phi, \quad N_{RC} = \varrho_{RC} \int n_{RC} d\Phi,$$

$$N_{NRM} = \varrho_{NRM} \int n_{NRM} d\Phi, \quad N_{NRC} = \varrho_{RM} \int n_{NRC} d\Phi,$$

$$C + G + \delta_s K_s + \xi \delta_e K_e = F(K_s, K_e, N_{NRC}, N_{NRM}, N_{RC}, N_{RM}).$$

3. Equations 17-21 hold.

¹⁶The time index is dropped from aggregate variables, given that this is characterization of the steady state.

4. The government budget balances:

$$G = \int \tau_k r(h + \Gamma) + \tau_c c + n\tau_l \left[\frac{nw(j, o, a, u)}{1 + \tilde{\tau}_{ss}} \right] d\Phi.$$

5. The social security system balances:

$$\int_{j \geq 45} \Psi d\Phi = \frac{\tilde{\tau}_{ss} + \tau_{ss}}{1 + \tilde{\tau}_{ss}} \left(\int_{j < 45} nw d\Phi \right).$$

6. The assets of the deceased at the beginning of the period are uniformly distributed among the living:

$$\Gamma \int \omega(j) d\Phi = \int (1 - \omega(j)) h d\Phi.$$

5 Estimating the production function

In this section, we describe the stochastic specification of the production function model, the equations to be estimated, and the results. The estimation strategy is as in [Krusell et al. \(2000\)](#). Parameter estimates from this section are used to produce a baseline calibration of the theoretical model, with which to run experiments of the impact of technological change on earnings dispersion. The data used in the estimation is described in [Appendix B](#).

5.1 Stochastic specification

The stochastic elements in our model are the unobserved technology components: (i) the relative technological level of the investment good sector; (ii) the set of labor-specific efficiency indices; and (iii) the factor-neutral technological process. We assume that the relative price of equipment ($\tilde{\xi}_t = \xi_t / \xi_{t-1}$) is trend stationary, and confirm this with a Dickey-Fuller test. We assume that the labor efficiency index processes have different linear trends for each labor variety. Defining the processes in logs we have:

$$\psi_t \equiv \log(\varrho_t), \quad \psi_t = \psi_0 + \psi_1 t + \nu_t, \quad (31)$$

where ψ_t is a (4×1) vector of the log of the latent efficiency indices, ψ_0 is a (4×1) vector of constants which specify the value of the indices at the beginning of the sample, ψ_1 is a (4×1) vector of growth rates, and ν_t is a (4×1) vector of shock processes that we

assume to be multivariate normal, i.i.d. with covariance matrix Ω : $v_t \sim N(0, \Omega)$. The i.i.d. assumption simplifies the identification of the factor-neutral technological change, A_t , which is described below.

5.2 Equation specification

We use a system with two sets of equations obtained from the first order conditions of agents in order to estimate the model: (i) the wage bills relative to the routine manual labor variety; and (ii) a no-arbitrage condition between investing in equipment and non-equipment capital. These are defined as follows:

$$\frac{w_{o,t} h_{o,t}}{w_{RM,t} h_{RM,t}} = wbr_{o,t}(\psi_t, X_t; \theta), \quad o \in O = \{\text{NRC}, \text{NRM}, \text{RC}\}, \quad (32)$$

and

$$1 + [F_{K_s}(\psi_{t+1}, X_{t+1}; \theta) - \delta_{s,t+1}] = E_t \left(\frac{\tilde{\zeta}_{t+1}}{\tilde{\zeta}_t} \right) (1 - \delta_{et+1}) + \frac{F_{K_e}(\psi_{t+1}, X_{t+1}; \theta)}{\tilde{\zeta}_t} \quad (33)$$

where 33 is obtained from equation 27, assuming that $\tilde{\zeta}_t \neq \tilde{\zeta}_{t+1}$, and where we substituted the return rates by factor marginal productivities.

Depreciation rates are indexed by t since they change over the time (see Appendix B). $wbr_{o,t}$ are functions of X_t and θ . X_t is the vector of inputs and depreciation rates $\{K_{s,t}, K_{e,t}, h_{\text{NRC},t}, h_{\text{NRM},t}, h_{\text{RC},t}, h_{\text{RM},t}, \delta_{s,t}, \delta_{e,t}\}$. The vector θ is the set of parameters $\{\alpha, \rho_1, \rho_2, \rho_3, \phi_1, \phi_2, \phi_3, \varphi_1, \varphi_2, \varphi_3, \psi_0, \psi_1, S, \eta_\omega, K_{e,0}\}$, including the first observation of the equipment capital stock, which we estimate jointly with the other parameters. η_ω is the standard deviation of the error term in the equipment price equation, which we specify below. Like Krusell et al. (2000), we assume that there is no risk premium in equation 33, and that tax treatment is identical between equipment and non-equipment capital returns. Finally, we substitute the first term on the right hand side of equation 33 with $E_t(\tilde{\zeta}_{t+1}/\tilde{\zeta}_t)(1 - \delta_{et}[1 - \tau_{k,t}]) + \omega_t$, where ω_t is the i.i.d. forecast error and $\omega_t \sim N(0, \eta_\omega^2)$. This set of assumptions imply that $A_t = Y_t/G(\cdot)$ from equation 15.

Data for the labor inputs in hours and the hourly (nominal) wages are used to obtain the left side of the set of equations 32. We use a measure of GDP at constant prices to find A_t .

The construction of the structure capital stock, depreciation rates, and relative prices is discussed on Appendix B. Given that this is a non-linear system of eight equations with unobserved state variables, standard linear Kalman filter techniques cannot be applied

to estimate the parameter vector θ . [Ohanian et al. \(1997\)](#) propose a two-step version of the SPML estimator to find θ for this type of problem, which we detail in Appendix C.

The parameter vector θ has dimension 36. Our sample contains 49 observations for each equation. We reduce the number of parameters to estimate by external calibration or by setting *a priori* restrictions. First, we impose that S be a diagonal matrix and that the variance of the disturbances be identical for all labor types. Thus, $S = \eta_v^2 I_4$, where η_v^2 is the common innovation variance and I_4 is a (4×4) identity matrix. Second, we fix $\psi_{4,0}$, the initial level of the latent efficiency index of routine manual workers, which is not identified. Third, we set the income share of structures to 0.04 as in [Krusell et al. \(2000\)](#). Finally, we regress the variation rate of the relative price of equipment on a linear trend in order to calibrate the forecast error variance of the equipment price index. We set η_ω to be equal to the estimated standard deviation of the error term in the regression $\tilde{\sigma}_\omega = 0.032$. This reduces the number of parameters to be estimated to 19: The common variance of the latent processes, η_v^2 , the elasticities, $\sigma, \rho_1, \rho_2, \rho_3$, the production function share parameters, $\phi_1, \phi_2, \phi_3, \varphi_1, \varphi_2, \varphi_3$, the parameters governing the latent state variables, except for $\psi_{4,0}$, and the initial level of capital equipment, $K_{e,0}$.

5.3 Estimation results and model fit

The model is estimated using data from 1967 to 2016 and the Simulated Pseudo Maximum Likelihood Estimation (SPMLE) procedure. Table 1 shows resulting estimates.

Elasticity estimates for the nested occupation types are all consistent with capital-occupation complementarity, i.e., $\sigma > \rho_i, i = 1, 2, 3$. The estimation of these elasticities is one of the contributions of this paper to the literature.

The most comparable estimates are provided by [Eden and Gaggl \(2018\)](#), who specify a CES production function with non-routine labor nested with capital. In contrast to our estimates of 0.5 and 2.1 for NRC and NRM labor, respectively, they estimate an elasticity of substitution of 1.4 for non-routine labor as a whole. For routine manual labor, their estimate is 8.0 for routine occupations, compared to our elasticity of 5.6 for RM. Although less comparable, [Krusell et al. \(2000\)](#) obtain a value of 0.67 for skilled labor, and 1.67 for unskilled labor. As to the processes of occupation-specific technology, we estimate that only the non-routine cognitive occupations have experienced positive growth, while routine manual labor has suffered the largest decline. We know of no other comparable estimates in the literature.

Figure 3 shows model fit to target moments over time. Figure 3a displays aggregate *ex post* return rates of equipment and structures implied by our model, which are zero

Table 1: Elasticity estimates


Parameter	Description	Value
σ	EOS RM	5.564
ρ_1	EOS NRC	0.497
ρ_2	EOS NRM	2.055
ρ_3	EOS RC	5.029
ϕ_1	Share NRC	0.378
ϕ_2	Share RM	0.086
ϕ_3	Share RM	0.279
φ_1	Share composite NRC	0.160
φ_2	Share composite NRM	0.045
φ_3	Share composite RC	0.023
$\psi_{0,1}$	Intercept NRC	0.859
$\psi_{0,2}$	Intercept NRM	1.936
$\psi_{0,3}$	Intercept RC	3.582
$\psi_{1,1}$	Slope NRC	0.002
$\psi_{1,2}$	Slope NRM	-0.006
$\psi_{1,3}$	Slope RC	-0.001
$\psi_{1,4}$	Slope RM	-0.010
$K_{e,0}$	Starting equipment capital	582

in expectation as per our assumption. They have a 4% average, as in [Krusell et al. \(2000\)](#), although a slightly increasing trend from the early 2000s onward.


Figure [3b](#) plots wage bill ratios implied by the model, as specified by the set of equations [\(32\)](#), and the data. Model predictions closely track the data. The NRC wage bill shoots up from close to on par with routine manual labor in 1968 to 3.5 in 2015. In contrast, NRM and RC wage bills grow slowly upwards relative to that of routine manual occupations, which is explained by both their lower level of complementarity with equipment capital as well as their declining level of latent efficiency.

Figure [3c](#) shows the model fit to the wage premia of each occupation relative to RM. As in the previous figure, the dashed lines indicate the data and the solid lines the model predictions. In all cases, the model tracks the data closely. This is important given that our goal is to use the estimated parameters to calibrate the theoretical model, and the key force driving earnings dispersion is the change in wage premia across groups.


Finally, Figure [3d](#) displays our estimate of total factor productivity in the U.S. for this period. From 1968 to 2008, TFP increased by almost 30% and then fell to around 20% in the following years. For comparison, the estimate of total factor productivity by the Penn World Table increases by 30% from 1968 to 2015 ([FRED](#)).




(a) Returns to capital



(b) Wage bill ratio by occupation



(c) Wage premia



(d) Total factor productivity

Note: The estimates presented cover the period from 1968 to 2015 as we lose both the first and the last period of the sample in order to estimate the model. In Figure 3d, total factor productivity is normalized to 1 in 1968. Construction of the measures is described in Appendix B.

Figure 3: Empirical model fit to target and non-target moments.

In conclusion, we provide new estimates for the elasticities of substitution between equipment capital and the occupation categories defined in Autor et al. (2003), which have been extensively used in the literature to discuss the impact of technological change and the future of labor markets. We find that our model is broadly compatible with the data, especially with respect to the occupation wage premia, which is crucial for ensuring that the predictions of the theoretical model are consistent with the data. We now turn to the calibration of the theoretical model, which uses the estimates obtained from this section to parameterize the production side of the economy.

6 Calibration

This section describes the calibration of the baseline model to match the U.S. economy in 1980. Parameters are either set directly (i.e., without solving the full model) to match their empirical counterparts, or estimated by simulated method of moments (SMM). Table 2 lists parameter values and sources.

6.1 External calibration

Preferences We set the inverse of the Frisch elasticity of labor supply, η , to 3 which is a standard value in the literature.

Labor productivity The wage profile through the life cycle (see equation 7) is calibrated directly from the data. We run the following regression, using Panel of Study of Income Dynamics (PSID) data:¹⁷

$$\ln(w_{it}) = a_i + \gamma_1 j + \gamma_2 j^2 + \gamma_3 j^3 + \varepsilon_{it}. \quad (34)$$

where j is the age of household i 's reference person and a_i is a household-specific effect. We then use the residuals of the equation to estimate the parameters governing the idiosyncratic shock, ρ and σ_ε . The scale parameters of the cost of choosing an occupation ($\mu_{\text{NRC}}, \mu_{\text{NRM}}, \mu_{\text{RC}}, \mu_{\text{RM}}$) are set such that they match the employment shares observed in 1980. The procedure is explained in Section 4. The location parameter, μ_{RM} , is normalized to 0.

Technology Equipment and structure depreciation rates are set to match those used in the estimation of the empirical model for 1980, and described in Appendix B. The production function is calibrated using the parameters estimated from the empirical model. The efficiency indices of each occupation are set to match those of the empirical model in 1980. The level of total factor productivity is set to the estimate from the empirical model for 1980.

Government We set θ_0 and θ_1 to the estimates obtained by Wu (2021) for 1980. For the social security rates we assume no progressivity. Both social security tax rates, employer and employee, are set to 0.06, the average rate in 1980. Finally, we set τ_c and τ_k to match the values obtained in Mendoza et al. (1994) for 1980, i.e, $\tau_c = 0.05$, $\tau_k = 0.47$.

¹⁷PSID data is described in section A of the Appendix.

Table 2: External calibration summary

Description	Parameter	Value	Source
Preferences			
Inverse Frisch elasticity	η	3.000	Assumption
Labor productivity			
Parameter 1 age profile of wages	γ_1	0.265	PSID
Parameter 2 age profile of wages	γ_2	-0.005	PSID
Parameter 3 age profile of wages	γ_3	0.000	PSID
Variance of idiosyncratic risk	σ_ϵ	0.307	PSID
Persistence idiosyncratic risk	ρ_u	0.335	PSID
Location of the cost of choosing NRC	μ_{NRC}	-5.712	CPS
Location of the cost of choosing NRM	μ_{NRM}	4.441	CPS
Location of the cost of choosing RC	μ_{RC}	0.379	CPS
Location of the cost of choosing RM	μ_{RM}	0.000	Assumption
Technology			
Equipment depreciation rate	δ_e	0.106	Section 5
Structures depreciation rate	δ_s	0.026	Section 5
Share structures	α	0.040	Section 5
Share NRC	ϕ_1	0.378	Section 5
Share NRM	ϕ_2	0.086	Section 5
Share RC	ϕ_3	0.279	Section 5
Share composite NRC	φ_1	0.160	Section 5
Share composite NRM	φ_2	0.045	Section 5
Share composite RC	φ_3	0.023	Section 5
EOS NRC	ρ_1	0.497	Section 5
EOS NRM	ρ_2	2.055	Section 5
EOS RC	ρ_3	5.029	Section 5
EOS RM	σ	5.564	Section 5
Latent efficiency NRC	ϱ_1	2.734	Section 5
Latent efficiency NRM	ϱ_2	4.955	Section 5
Latent efficiency RC	ϱ_3	34.662	Section 5
Latent efficiency RM	ϱ_4	0.378	Section 5
Total factor productivity	A	16.728	Section 5
Relative price of investment goods	ξ	1.000	Assumption
Government and SS			
Consumption tax rate	τ_c	0.054	Mendoza et al. (1994)
Capital income tax rate	τ_k	0.469	Mendoza et al. (1994)
Tax scale parameter	θ_0	0.850	Wu (2020)
Tax progressivity parameter	θ_1	0.187	Wu (2020)
SS tax employees	τ_{ss}	0.061	Social Security Bulletin, July 1981
SS tax employers	$\tilde{\tau}_{ss}$	0.061	Social Security Bulletin, July 1981

6.2 Internal calibration

To calibrate the parameters for which we do not have a direct empirical counterparts, $\{\beta, \chi, \varphi, \sigma_{NRC}, \sigma_{NRM}, \sigma_{RC}, \sigma_{RM}\}$, we use a simulated method of moments, for which we construct the following loss function:

$$L(\tilde{\theta}) = ||M_m - M_d||, \quad (35)$$

where $\tilde{\theta}$ is the vector of parameters to be estimated and M_m and M_d the moments in the model and in 1980, respectively. Our estimate, $\tilde{\theta}^*$, is obtained by minimizing (35).

We use the ratio between average wealth of 65 and older to the average wealth in the economy as the target for the utility of bequests parameter. The discount factor is set by targeting the capital-to-output ratio. The capital stock is obtained from the estimation of the empirical model of section 5. Disutility from work targets hours worked, and we calibrate the occupation-specific variances of ability to target the variance of log earnings observed in the data for each occupation. Calibration fit is presented on Table 3. Table 4 presents the parameters calibrated internally.

Table 3: Calibration fit

Data moment	Description	Source	Model	Data
65-on/all	Average wealth of households 65 and over	US Census Bureau	1.310	1.311
K/Y	Capital to output	BEA and CPS	1.412	1.412
\bar{n}	Fraction of hours worked	BEA	1/3	1/3
$\text{Var ln}(e_{NRC})$	Variance of log earnings (NRC)	CPS	0.408	0.409
$\text{Var ln}(e_{NRM})$	Variance of log earnings (NRM)	CPS	0.410	0.406
$\text{Var ln}(e_{RC})$	Variance of log earnings (RC)	CPS	0.409	0.410
$\text{Var ln}(e_{RM})$	Variance of log earnings (RM)	CPS	0.305	0.304

Table 4: Parameters Calibrated Internally

Parameter	Value	Description
φ	9.993	Bequest utility
β	0.961	Discount factor
χ	66.981	Disutility of work
$\sigma_{a,NRC}$	0.519	Variance of ability NRC
$\sigma_{a,NRM}$	0.515	Variance of ability NRM
$\sigma_{a,RC}$	0.517	Variance of ability RC
$\sigma_{a,RM}$	0.385	Variance of ability RM

7 Quantitative Results

In this section, we use our model calibrated to the U.S. economy in 1980 in order to answer the two main questions described in the introduction: To what extent does technological change explain the observed increase in earnings dispersion? How does technological change affect the optimal progressivity of the tax system?

7.1 Sources of variation in earnings inequality

The main experiment conducted in this section is to recalibrate the model to match technological and tax changes from 1980 to 2015. We then decompose the variation in the earnings inequality statistics between steady states to identify the role of investment-specific technological change, occupation-specific technological change, and TFP growth on the variation of earnings dispersion.

We also compare the magnitude of these sources of variation in earnings dispersion to others, such as the observed changes in the progressivity of the tax system. The decompositions are carried out by setting the relevant parameters to their 1980 level and comparing the resulting change in earnings dispersion to its 2015 value as calculated in the baseline experiment.

Parameters related to tastes, individual productivity processes and the production function are kept constant between steady states: The age profile of wages ($\gamma_1, \gamma_2, \gamma_3$), the idiosyncratic productivity process (ρ_u and σ_ϵ), preferences (λ, η, β), ability variance parameters ($\sigma_{a,o}, o \in O$), and production function shares and elasticities. Parameters changed from 1980 to 2015 are listed on Table 5. The main shifts are in technology parameters, idiosyncratic occupation costs, and tax rates.

In the new steady state, we set the relative price of investment goods to 40% of the initial price index, which mimics the fall measured in the data between 1980 and 2015. Labor efficiency indices are set to their 2015 levels, using the functional forms of those processes estimated in section 5. Likewise, TFP is set to equal the estimated level in 2015. The location parameters of the idiosyncratic cost distributions are set such that they match the occupation employment shares observed in 2015.

The scale and the progressivity parameters of the labor income tax schedule are set to match the estimate of Wu (2021). Social Security tax rates are those described in Brinca et al. (2016) for the U.S. economy. Both the consumption tax and the capital income tax are calculated using the method in Mendoza et al. (1994).

Table 6 contains the fit of the model moments to the data in 1980 and 2015. In the first section of the table, we compare relative input quantities from the theoretical model to

Table 5: Parameter Changes

Parameter	Description	1980	New SS
τ_c	Consumption tax	0.054	0.050
τ_k	Capital income tax	0.469	0.360
τ_{ss}	Employee SS tax	0.061	0.077
$\tilde{\tau}_{ss}$	Employer SS tax	0.061	0.077
θ_0	Tax scale	0.850	0.922
θ_1	Tax progressivity	0.187	0.137
ξ	Investment price	1.000	0.405
ϱ_1	Latent efficiency NRC	2.734	2.986
ϱ_2	Latent efficiency NRM	4.955	4.051
ϱ_3	Latent efficiency RC	34.662	33.907
ϱ_4	Latent efficiency RM	0.378	0.267
μ_{NRC}	NRC cost location parameter	-5.713	-7.618
μ_{NRM}	NRM cost location parameter	4.441	3.938
μ_{RC}	RC cost location parameter	0.379	-2.332

those obtained from estimating the empirical model of the production function in section 5. Relative inputs quantities are fairly close to our estimates, with the exception of the growth of equipment capital, which the theoretical model substantially underestimates relative to the empirical model. In other words, the model is unable to generate a rise in savings comparable to our empirical estimates.¹⁸

The second section shows the wage changes by occupation between both steady states. The model overestimates wage growth for all occupations, with the exception of NRM. However, as can be seen in the next section of the table, wage premia levels and changes are very close to the data in both years, which is key in terms of accounting for the change in earnings dispersion.

The third section displays measures of wage inequality. In 1980, the model predicts an overall level of wage dispersion which is above its data counterpart, except for the RM occupation. Thanks to the mechanisms which widen between-group inequality across time, this gap is reduced in 2015 for the total variance of log wages, but not for within-occupation wage dispersion, which remains at its 1980 levels due to the calibration strategy implemented. This strongly suggests that there are other forces in the U.S. economy driving within-occupation wage inequality which are not present in the model, such as increased women's labor force participation and higher human capital accumu-

¹⁸This is likely due to the fact that the U.S. is a large open economy. This means that the stock of capital has grown not only via increased domestic savings but also through foreign direct investment (FDI). According to the BEA, the stock of FDI in the U.S. increased 40 fold between 1980 and 2015 from 83\$ billion to 3.4\$ trillion (Source: BEA annual data on FDI position).

Table 6: Theoretical model fit

Variable	1980		2015	
	Model	Data	Model	Data
K_e / N_{NRC}	6.27	7.80	16.70	39.14
$N_{\text{NRM}} / N_{\text{NRC}}$	0.66	0.55	0.49	0.43
$N_{\text{RC}} / N_{\text{NRC}}$	10.61	9.10	5.74	5.66
$N_{\text{RM}} / N_{\text{NRC}}$	0.16	0.16	0.05	0.05
NRC wage	1.00	1.00	1.38	1.28
NRM wage	1.00	1.00	1.11	1.11
RC wage	1.00	1.00	1.25	1.14
RM wage	1.00	1.00	1.00	0.93
<i>Wage premia</i>				
NRC	1.35	1.31	1.83	1.80
NRM	0.60	0.63	0.66	0.74
RC	0.91	0.88	1.14	1.09
<i>Variance of log wages</i>				
NRC	0.38	0.31	0.48	0.44
NRM	0.36	0.29	0.36	0.41
RC	0.36	0.21	0.36	0.25
RC	0.36	0.25	0.36	0.36
RM	0.26	0.26	0.27	0.33
<i>Variance of log earnings</i>				
NRC	0.43	0.45	0.57	0.57
NRM	0.41	0.41	0.43	0.51
NRM	0.41	0.41	0.45	0.37
RC	0.41	0.41	0.44	0.49
RM	0.30	0.30	0.33	0.39

Note: The first section of the table indicates the relative input quantities. The second section indicates wage per efficiency unit by occupation (model definition) and wages per hour at constant 1968 prices by occupation (data definition). Both prices are normalized to 1 in 1980. The third section indicates the wage premia calculated as the ratio between the marginal productivities of the labor from each occupation category relative to RM in the case of the model. The empirical counterpart of the model wage premia is described in Appendix B.

lation, for example. Nonetheless, of the 42% increase in the variance of log wages the model is able to replicate 60%.

The final section shows the variance of log earnings which is the centrepiece of our analysis. Total variance of log earnings grows 33% from 1980 to 2015 in the data and 27%

in the model, implying that the model overshoots the growth in earnings inequality. The change in earnings dispersion results from two sources: Within and between-occupation variance. While the model matches the within-occupation earnings dispersion in 1980, the 2015 value is non-targeted to allow the mechanisms of the model to endogenously generate a prediction which can be compared to the data. As a result, within-occupation earnings dispersion in the model rises only slightly between 1980 and 2015. In contrast, it rises significantly for most occupation types in the data. In particular for the NRC which form a growing fraction of total employment (30% in 1980 and 41% in 2015).

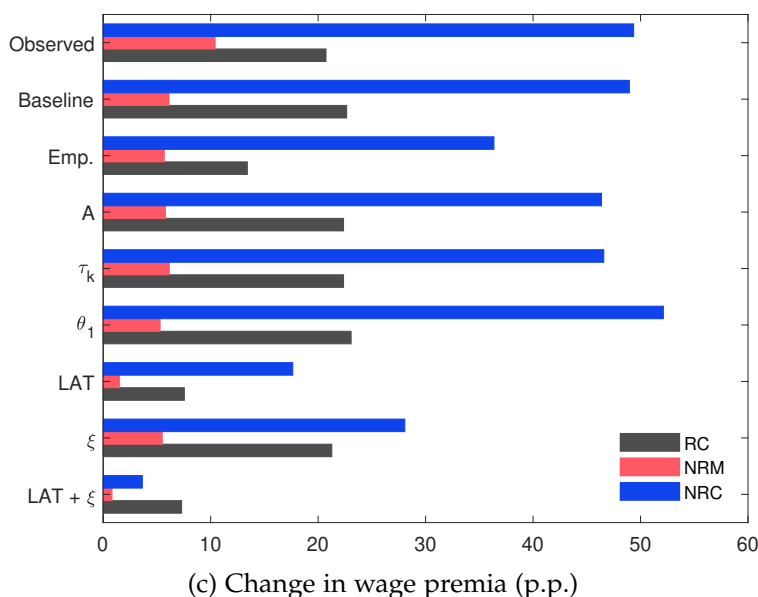
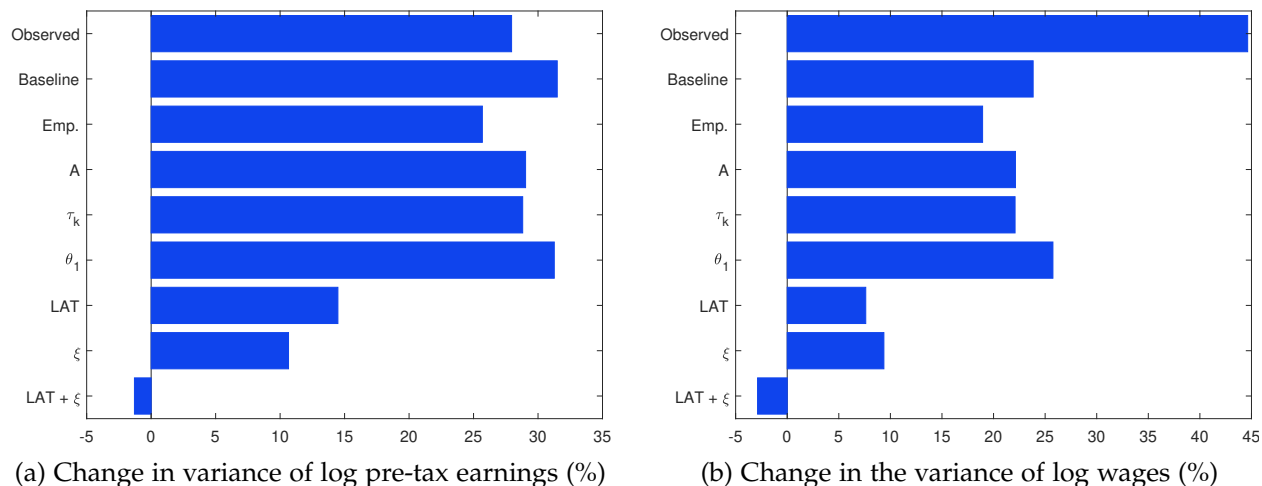
To understand the drivers of the change in aggregate earnings inequality, we decompose the model prediction of the variation in labor market dispersion measures. This is achieved by setting each set of parameters of interest to their 1980 levels while keeping the remaining parameters at their 2015 level. We then compare the resulting change to the variation produced by the baseline experiment.

Figure 4 illustrates these exercises by displaying the response of labor income dispersion measures to the set of parameter shifts presented on Table 5. The first bar in each panel indicates the observed change in that measure, while the second bar indicates the change predicted by the model as a result of the baseline parameter shifts shown in Table 5.

Figure 4a, in the top left panel, shows how the model fares in generating a shift in pre-tax earnings inequality comparable to the one observed in the data. As previously mentioned, the model prediction slightly overshoots the increase in the variation of earnings inequality observed in the data (33% in the model but only 27% in the data). However, this confirms that our framework is successful in producing a prediction regarding labor market inequality which is comparable in magnitude to the one observed in the data.

The drop in the relative price of investment, ζ , is the single most relevant source of the increase in pre-tax earnings inequality. If this drop had not taken place, our framework predicts a rise of only 10%. Thus, investment-specific technological change alone accounts for two thirds of the model-predicted variation. This follows from the increased dispersion of between-group wages, resulting from the complementarity effect between occupations and equipment capital.

Occupation-biased technological change is next in terms of relevance: Keeping the occupation-specific efficiency indices at their 1980 levels generates only a 15% increase in pre-tax earnings dispersion, half of the total change predicted by the model. This results from lower surge in wage variance (only 20% of the baseline variation) and in the NRC wage premium in particular, which rises by only 18 p.p.



Note: The bar denoted “Observed” indicates the change in the indicator recorded in the data between 1980 and 2015. The data used is from the CPS and is described in section 3. “Baseline” indicates the change predicted in the theoretical model from 1980 to 2015. Each of the remaining bars indicate the change in the model statistics resulting from keeping the corresponding parameters at their 1980 levels. “Emp” represents the impact of keeping the location parameter of the distributions of idiosyncratic costs of entering a given occupation at their 1980. “A” is total factor productivity. “LAT” is the set of occupation-specific efficiency indices.

Figure 4: Decomposition of the variation in labor income dispersion measures from 1980 to 2015. Technological progress accounts fully accounts for the changes in earnings inequality predicted by the model.

Taken together, ISTC and LAT fully account for the increase in pre-tax earnings dispersion. Keeping both sets of parameters at their 1980 values yields a reduction in wage variance and eliminates the role of the NRC wage premium in driving the change

in inequality.

In contrast, other sources of variation in the variance of pre-tax log-earnings in are much less relevant. The reduction in progressivity had a positive but comparatively much smaller effect. Changes in the costs of acquiring the skills necessary to join each occupation, while having impact on the changes in employment shares (Figure 5) produced only a small increase in earnings inequality.


Figure 4b, on the top right, compares the observed change in the variance of log wages to the baseline model change. In contrast to the previous measure, the increase in the variance of log wages is about half of what is observed in the data. Part of the reason for this observation is the fact that the variance of log wages is non-targeted in 1980, and is above what is recorded in the data, so the starting point is larger in the model. However, 4b shows once again that technological change is a major driver of wage inequality, given that the variance of log wages would actually drop between 1980 and 2015 in the absence of ISTC and labor-augmenting technological change.

Figures 4c and 5 show how relative prices and quantities of labor by occupation respond to the experiments conducted. This enables us to understand the direction of mechanisms and their strength.


With respect to prices, there are two competing forces: On the one hand, the increase in the NRM wage premium reduces earnings inequality, all else equal, given that it starts out in negative territory in 1980 and increases 6 p.p. (10 p.p. in the data). In other words, technological narrows the gap between the wage of NRM occupations, which lie at the bottom of the wage distribution, and RM occupations. On the other hand, the wage premia of NRC and RC occupations increase relative to RM occupations. In particular, the NRC wage premium increases 50 p.p., and stands as one of the main sources of increased earnings dispersion in this model.

With respect to employment shares, the main takeaway is that technological change generates an increases in the weight of non-routine occupations in total employment. Without occupation-biased technological change and ISTC, the model predicts that the share of NRC workers would actually drop in 2015 with respect to 1980, and the share of NRM would remain unchanged.


Finally, we could be underestimating an important source of increasing earnings dispersion: The large and increasing weight of NRC occupations in total employment. As the results on table 6 indicate, the variance of log wages by occupation increased substantially, especially for NRC occupations. The model underpredicts this increase, which could drive up earnings inequality still further as a larger proportion of the workforce join a set of professions with greater wage variance.




(a) NRC employment share (p.p.)



(b) NRM employment share (p.p.)



(c) RC employment share (p.p.)



(d) RM employment share (p.p.)

Note: See Figure 4 for the description of the y-axis.

Figure 5: Decomposition of variation in employment shares by occupation type from 1980 to 2015.

In summary, technological change (and ISTC in particular) is able to generate an increase in pre-tax earnings inequality which is comparable to the one observed in the data. At first glance, this would imply a strengthening of the case for an increase in the progressivity of the labor income tax system.

In contrast, according to Wu (2021), the U.S. tax system was less progressive in 2015 than in 1980. Moreover, we find that optimal progressivity dropped during that period. In the next two subsections we show optimal progressivity in 1980 and 2015 and discuss how technological change has affected it.

7.2 Optimal progressivity in 1980

In this subsection we examine the model-implied welfare function with respect to the progressivity of the labor income tax system, θ_1 , and decompose it into different welfare effects (à la Flodén, 2001) and occupations. Our main finding is that optimal progressivity is lower than the estimated progressivity of the U.S. labor income tax system in 1980. Nonetheless, the model-implied gains from moving to the optimal level of θ_1 are very low at 0.06% in consumption equivalent variation.


To make these experiments compatible with the literature on optimal taxation we use θ_0 , the scale parameter of distortionary labor income taxation, as a clearing variable for the government budget constraint. In other words, θ_0 is such that the government is able to raise enough tax revenue to cover the level of government expenditure, G , in the initial steady state.

To measure consumer welfare, we use the *utilitarian social welfare* criterion (i.e., the expected lifetime utility of a household who is yet to enter the labor market).¹⁹ In a nutshell, the welfare gain from choosing a progressivity level $\theta_1^B \neq \theta_1$ can be broken down into three elements: (i) the gain from reducing *uncertainty* that agents face, (ii) the gain from reducing *inequality* in average lifetime marginal utilities of consumption and leisure, and (iii) the impact that progressivity has on the overall *levels* of consumption and leisure via the incentive to work and invest. Figure 6 shows the results of our analysis.

The left panel of Figure 6 plots the social welfare function for 1980 (the solid blue line). The model-implied optimal progressivity is 0.15, below the estimate of actual progressivity in 1980 of 0.19 by Wu (2021). Our estimate of optimal progressivity for 1980 is close but below that of Heathcote et al. (2020), who put it at 0.18. Like them, we focus only on comparative statics and do not take into account the transition path to the new steady state. Unlike Heathcote et al. (2020), we do not have cross-sectional variation in the disutility of labor, for example, which may partly explain the difference in results.

The low welfare gain from moving to $\theta_1 = 0.15$ masks large countervailing movements in its components which form the basis of the trade-offs faced by the policy-maker. On the one hand, a reduction θ_1 hampers the effectiveness of the labor income tax system as a mechanism to reduce both the uncertainty faced by agents and inequality. On the other hand, it increases the incentive to exert effort, to save, and affects prices and occupational choice. Figure 7 shows the comparative statics of output, wages and employment shares with respect to progressivity in 1980. We look at the positive and

¹⁹See Appendix E for a definition of the welfare criterion and an exposition of how to decompose it into different effects in the case of our framework.



Note: The left panel plots social welfare as a function of the progressivity parameter, θ_1 , for the 1980 calibration. Social welfare is measured as the consumption equivalent variation required for agents entering the labor market to be indifferent between the baseline policy and the new one, without accounting for transitions. The vertical lines indicate current and optimal progressivity levels. θ_1^* indicates optimal progressivity according to the model. The right panel shows the social welfare functions for households starting their life in the indicated occupation categories in each year. In this case, welfare is the expected lifetime utility of agents once they have chosen an occupation but before they know a_i (ability) and ϵ_{i1} (starting level of wage risk).


Figure 6: Optimal progressivity in 1980 and across occupations. Optimal progressivity is lower than estimate for 1980.

negative effects of a reduction progressivity in turn.


On the plus side, reducing θ_1 to 0.15 increases aggregate saving and effort by reducing the marginal tax rate on top wage earners. This raises the level of output per capita by 2% (Figure 7a). The net effect on wages depends on two factors: (i) the direct effect of increased saving and capital on the marginal productivity of each occupation, and (ii) on how agents will change their occupational choice.

Figures 7c and 7d show these channels at work in equilibrium: Reducing progressivity increases the attractiveness of higher paid occupations which, everything else equal, increases selection into these occupations. In our case, highest paid workers are, on average, in non-routine cognitive occupations. Therefore, the share of employment in NRC occupations expands as θ_1 contracts. This inflow of employment into NRC raises the marginal productivities of other occupations and, thus, has a positive effect on their wages. Reducing θ_1 to its optimal value implies a 0.3% drop in NRC wages in equilibrium.


However, the interests of individuals in a given occupation do not necessarily co-




(a) Capital and the interest rate



(b) Output and hours



(c) Wages by occupation



(d) Employment shares by occupation

Note: Total labor input is the sum of labor efficiency units supplied in the economy. Average hours is the percentage of the labor endowment used to work on average. Each panel shows how prices and quantities in the economy change with respect to progressivity.

Figure 7: Comparative statics with respect to progressivity in 1980.

incide. This is most obvious for agents in the middle or at the top of the NRC wage distribution. These individuals benefit from an increase in progressivity, given that it discourages agents at the bottom of the wage distribution from joining their occupation, reduces NRC labor input, and raises their marginal productivity and wages. This is the mechanism that explains why, in Figure 6b, agents who chose NRC occupations prefer, on average, a high level of progressivity despite being at the top of the wage distribution.²⁰

The same (reverse) logic applies to agents in the occupations at the bottom of the wage distribution. While they may benefit from an increase in the progressivity of the

²⁰This is the same mechanism that underlies the creation of professional guilds.

tax system, it also increases the attractiveness of joining their occupation, all else equal. This is the case of non-routine manual occupations, which earn the lowest wages in the economy, on average. A drop in progressivity to the optimal value increases wages by 1.5% and reduces their employment share by 0.4%. The effect on NRM wages is so strong that it actually starts to revert the net welfare loss for very low levels of progressivity, as scarcity of NRM workers implies that their wage skyrocket.

For middle wage earners, in RC and RM occupations, wages increase by 0.6% and 0.3% and employment shares drop by 0.3 p.p. and 0.4 p.p., respectively. In the end, increased output and wages for the lower-paid occupations dominate.

On the minus side, lowering progressivity reduces the interest rate through greater capital accumulation. This hampers the ability of agents to insure against risk and raises the cost of facing uncertainty. It also reduces the ability of the tax system to reduce inequality.

In the end, the positive effects on output and wages from setting $\theta_1 = 0.15$ exceed the negative effects from greater inequality and lower insurance against uncertainty, but just barely. At the heart of the trade-off that determines optimal progressivity lies occupational choice. Whereas in most of the previous literature it only affected the intensive margin of the labor choice (either through hours worked or continuous skill choice), it now affects the extensive margin via the choice of occupation, which makes top-wage earners ambivalent as to their preference of progressivity of the tax system.

In the next subsection, we see how these forces are exacerbated by technological change.

7.3 The impact of technological change on optimal tax progressivity

In this section we ask the second of the main questions in our paper: How did the technological transformation that took place between 1980 and 2015 affect optimal progressivity? To answer this, we evaluate optimal progressivity in 2015 and then investigate how each source of technological change affects our answer.

In a nutshell, our results show that technological change had a major impact in reducing optimal progressivity. In particular, ISTC was the most relevant factor for this outcome. We begin by discussing the main implications of each source of technological change on aggregate variables. Then, we conduct a welfare analysis for 2015 and disentangle the drivers of our conclusion. Table 5 shows the parameter changes that mimic the move to the 2015 calibration. Table 7 shows how removing each source of technological change affects output per capita, capital, and prices.

Table 7: Impact of technological change on quantities and prices

	2015	No ISTC	1980 LAT	1980 TFP	No tech Δ
Output p/ capita	1.39	1.08	1.35	1.23	0.95
Capital to output	4.96	0.90	4.57	4.53	0.78
Interest rate (post-tax)	0.05	0.03	0.05	0.05	0.03
NRC wage	1.38	1.04	1.36	1.22	0.95
NRM wage	1.11	0.93	1.23	0.99	0.94
RC wage	1.25	1.05	1.30	1.12	1.00
RM wage	1.00	0.85	1.20	0.90	0.93

Note: Table shows the equilibrium impact of technological change on quantities and prices in the model. “No ISTC”, “1980 LAT”, “1980 TFP”, and “No tech” denote the values of aggregate variables by, respectively, setting investment prices at the 1980 value, the occupation-specific efficiency indices at 1980 values, TFP at its 1980 level, and removing all sources of technological change while keeping other parameters at their 2015 values. All variable (except the interest rate) are equal to 1 in 1980.

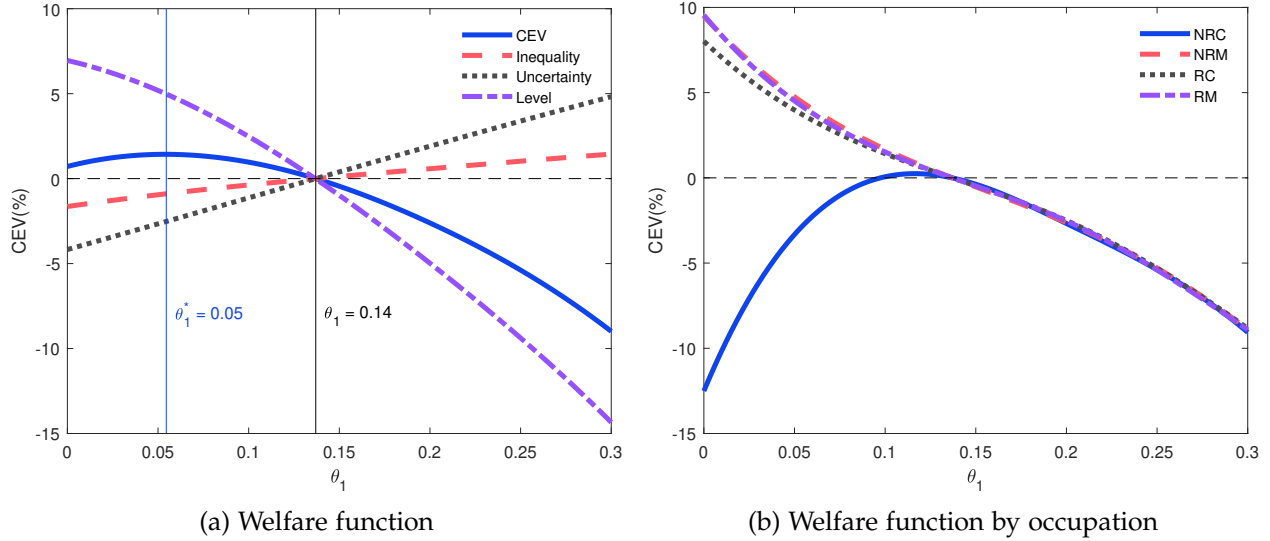
The main takeaway from Table 7 is that ISTC had the most significant impact on aggregate variables out of all sources of technological change. First, it accounts for around 70% of output per capita growth compared to a scenario where there is no tech growth between 1980 and 2015, via making investment in equipment capital more attractive. This is reflected in a higher post-tax return rate when compared with all other scenarios. Second, it is responsible for between 80 to 100% of the growth in the hourly wages in all occupations compared with the scenario where there is no technological change.

Labor-augmenting technological change has a positive but very small impact on GDP per capita and capital. However, it has a negative impact on the wages of all occupations except for NRC. This follows from the results of the production function estimation.

Finally, TFP has a large positive impact on GDP per capita, accounting for about 30% of its growth.²¹ It is responsible for between 40 to 100% of the wage growth by occupation relative to the no tech growth scenario.

In summary, technological change produced a significant increase in the level of output per capita, interest rates, and wages. However, as we argued in section 7.1 it also generated an increase in labor market inequality, with the NRC wages growing more than those of other occupations. Clearly, this affects the trade-off between efficiency and inequality that underlies the determination of optimal progressivity. Figure 8 displays the welfare analysis.

²¹The decompositions need not add to 100%, as they interact.



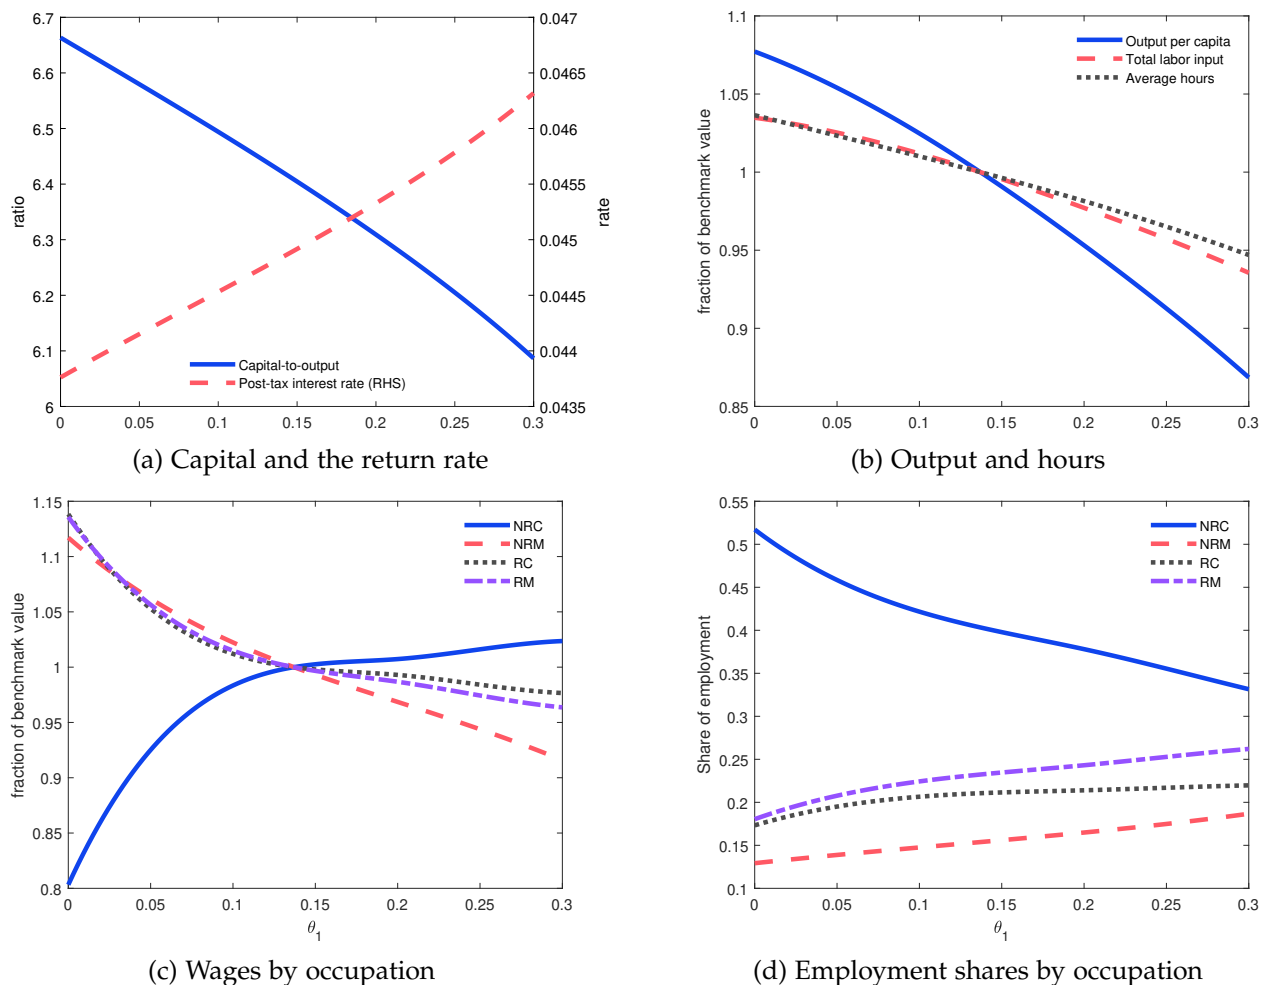
Note: The top left panel plots social welfare as a function of the progressivity parameter, θ_1 , for 2015, under the baseline calibration. Social welfare is measured as the consumption equivalent variation required for agents entering the labor market to be indifferent between the baseline policy and the new one, without accounting for transitions. The vertical lines indicate current and optimal progressivity levels. θ_1^* indicates optimal progressivity according to the model. The right panel shows the social welfare functions for households starting their life in the indicated occupation categories in each year. In this case, welfare is the expected lifetime utility of agents once they have chosen an occupation but before they know a_i (ability) and ϵ_{i1} (starting level of wage risk).

Figure 8: Optimal progressivity in 2015 and across occupations. Large reduction in optimal progressivity.

Wu (2021) estimates that the actual progressivity in the U.S. tax system declined from 0.19 to 0.14 between 1980 and 2015. In our framework we predict that optimal progressivity dropped from 0.15 in 1980 to 0.05 in 2015. If the economy were to instantly move to the new policy steady state, unborn agents would be better off by about 1.4% in consumption equivalent variation.

What lead to this remarkable two-thirds drop in optimal progressivity? What sources of technological transformation explain this shift in optimal progressivity? How do they affect the trade-offs faced by the policymaker? To answer these questions, we look at the decomposition of the welfare function and disentangle the effects of each source of technology (Figures 10 and 11). In a nutshell, we find that ISTC plays the key role in driving the results.

The trade-off between efficiency and redistribution is now more significantly dominated by the efficiency side (Figures 8a and 11). This is due to three main reasons: (i) the higher proportion of NRC workers implies that increasing capital raises output by a greater amount, given their high complementarity with equipment; (ii) the effect



Note: Total labor input is the sum of labor efficiency units supplied in the economy. Average hours is the percentage of the labor endowment used to work on average. Panels show how prices and quantities in the economy change with respect to progressivity.

Figure 9: Comparative statics of progressivity in 2015.

of capital accumulation during the 1980-2015 period on the sensitivity of marginal productivities and employment shares, and (ii) the higher equilibrium return rate on capital given its high productivity. The latter implies that the welfare gains from improving risk-sharing via more progressive taxation are less significant, which reduces the importance of the *uncertainty* channel relative to 1980. The effects described in (ii) are ambiguous and must be looked at in turn.

On the one hand, because of capital-occupation complementarity the wages of different occupations will rise by different amounts in response to an increase in capital equipment due to larger savings, leading to rising dispersion in marginal utilities of consumption and leisure, which negatively affect the *inequality* channel.

On the other hand, employment shares change in equilibrium as agents select different occupations in response to variation in relative wages and post-tax earnings. For example, a reduction in θ_1 will make high wage earners better-off by increasing after-tax earnings. However, it will have a negative effect on their wages by making their occupation more attractive to work in via both the intensive and the extensive margins. Figure 13 shows how quantities and prices respond to θ_1 in equilibrium. Our framework predicts that the net effect of reducing θ_1 to 0.05 is to significantly increase the wages of all occupations (NRM: 5.8%; RC: 4.8%; RM: 5.1%) except for NRC (-6.7%). The NRC occupation experiences an increase in its share of total employment (5.1 p.p.), while all others experience a drop (NRM: -1.4 p.p.; RC: -1.4 p.p.; RM: -2.3 p.p.).


In conclusion, the impact of a reduction of progressivity in the after-tax earnings of NRC occupations, which are the highest-paying, makes them more attractive and offsets the occupation-capital complementarity effect embedded in the production function.

The relative strength of these competing mechanisms explains the preferences for optimal progressivity by occupation explicit in Figure 13. Are the agents in NRC occupations partial to a reduction in progressivity? Strikingly, the answer is that they are less keen on that policy than the workers in the lower-paid occupations. The reason is the effect that a lower progressivity policy has on occupation selection: Lowering the marginal tax rate on high wage earners increases their after-tax earnings, but leads to an inflow of workers and a reduction in their hourly wages at some point.²² As progressivity drops, so does the share of NRC workers in the economy increase, as well as their weight in determining the optimal redistributive policy.

In contrast, lower paid occupations would unanimously benefit from a flat tax rate. On the one hand, there is no redistribution for agents to face uncertainty. However, as previously discussed, the higher return rate on savings relative to 1980 change dampens this channel, given that insurance is cheaper. The inequality channel is also negatively affected, but has a low weight, given that a large fraction of wage inequality in this economy is endogenous and depends on individuals optimally choosing their occupation. On the other hand, a flat tax rate leads to a surge in savings, in the capital-to-output ratio, and an outflow of workers to higher-paid occupations. Both of these mechanisms have the unambiguous impact of increasing the wages of all occupations except for the NRC, whose members cannot move to a higher-paid occupation.

These results stand in stark contrast with those by [Heathcote et al. \(2010\)](#), who find an optimal level of $\theta_1 = 0.16$ in 2016. Part of the reason for this is the heterogeneous impact of savings and technological change on the marginal productivities of each occupation,

²²That point being $\theta_1 = 0.11$.



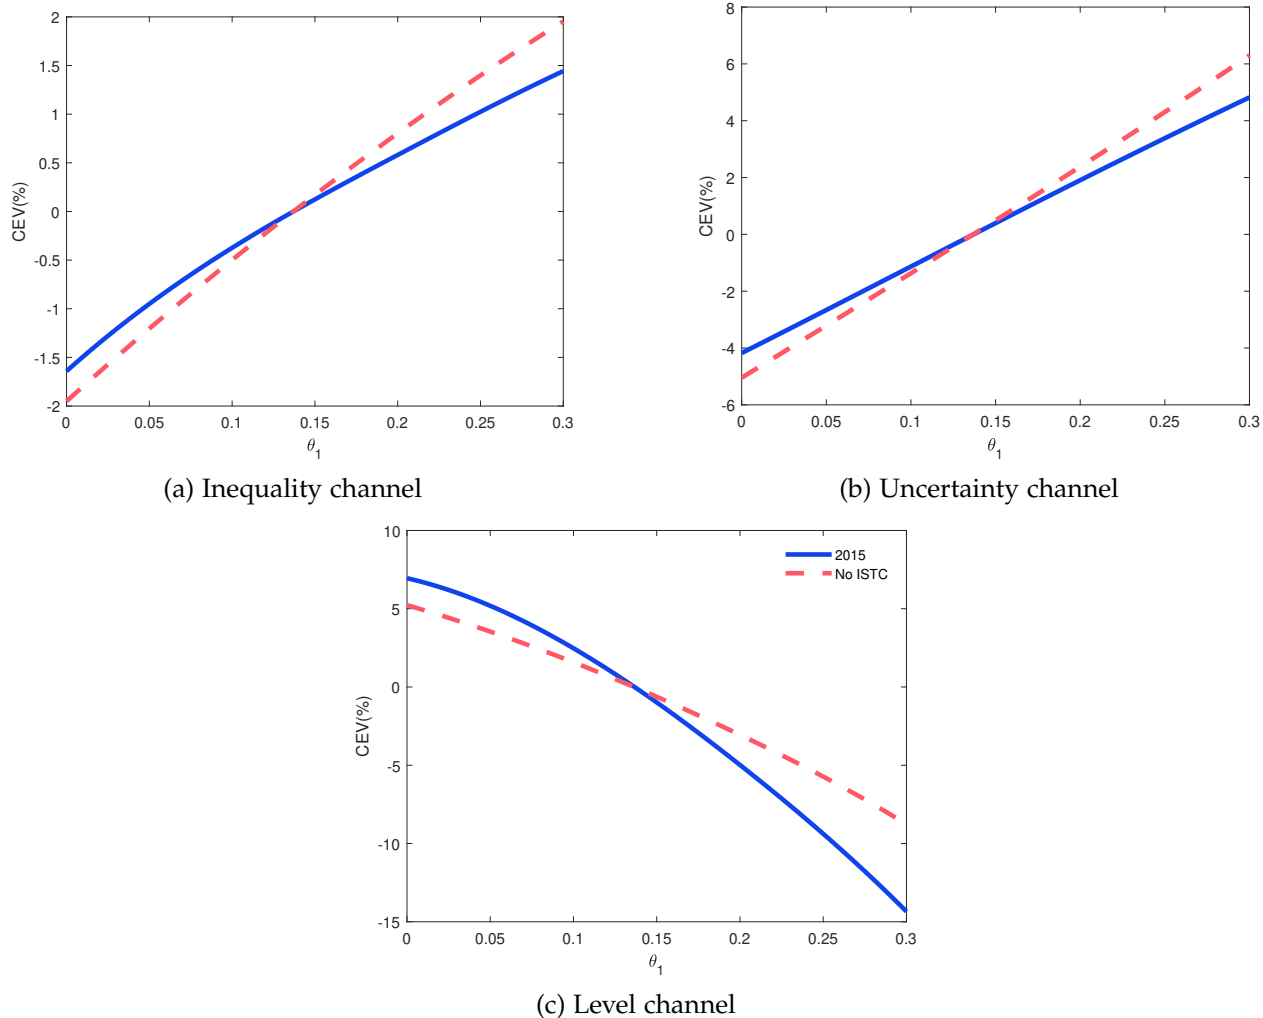
Note: The solid blue line denoted “2015” indicates the social welfare function for the economy calibrated to 2015. The dotted red line denoted “No ISTC” is the welfare function of the economy calibrated to 2015 but where the relative price of investment goods is kept at its 1980 level. The dotted purple line denoted “1980 LAT” is the welfare function for the economy calibrated to 2015 but where occupation-specific efficiency indices have been kept at their 1980 levels. The grey dash-dotted line denoted “1980 TFP” shows the welfare function with TFP at 1980 levels. θ^* indicates optimal progressivity according to the model for each experiment.

Figure 10: Optimal progressivity in 2015 and the role of technological change. ISTC fully accounts for the leftward shift in the welfare function estimated in 2015.

which is absent from their model. Rather, in their framework agents choose their level of skills, which are imperfect substitutes in production and have a single constant elasticity of substitution, and there is no role for capital in production.

Figure 10 confirms our interpretation of ISTC as the main driver of these changes. The solid blue line is the welfare function in 2015, where the maximizer is 0.05. The dashed red line indicates the calibration for 2015 with equipment investment prices at their 1980 level. Here, the maximizer is 0.18, close to the value estimated for 1980 and the optimal θ_1 estimated by Heathcote et al. (2019). As discussed earlier, this equilibrium displays lower return rates, and lower sensitivity of wages, employment shares, and output per capita to progressivity. As a result, the optimal progressivity is much higher.

The purple dotted line indicates the welfare function in an economy calibrated to 2015 with occupation-specific efficiency indices at their 1980 levels. In this case, earn-



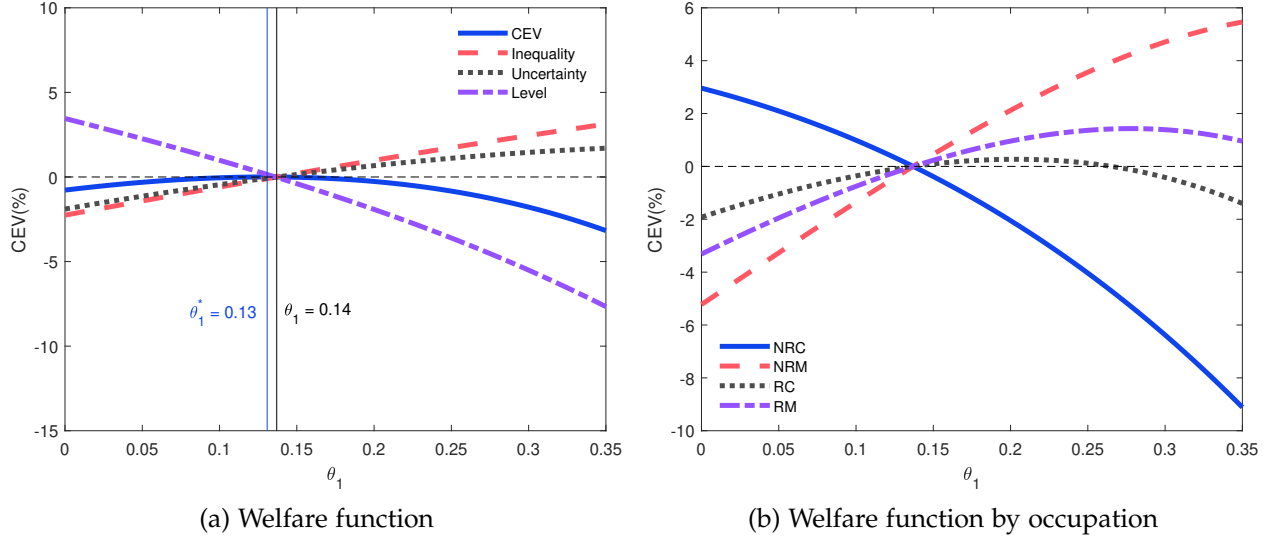
Note: Each graph shows the impact of changing progressivity on each of the elements composing the welfare function. “2015” indicates the breakdown for the 2015 calibration. “No ISTC” shows the breakdown for the 2015 calibration in the absence of investment-specific technological change.

Figure 11: Welfare trade-offs with and without ISTC.

ings inequality is lower, given that the NRC occupation does not experience as large an increase in wages as in the 2015 calibration. As a result, the inequality channel is even weaker, which drives optimal progressivity further to the left.

Finally, the dash-dotted grey line plots the welfare function when keeping TFP at its 1980 value. As mentioned before, TFP raises wages and output, but has little effect on changing the trade-offs faced by the policymaker in 2015.

Figure 11 confirms our hypothesis that ISTC increases the sensitivity of the level effect to progressivity and weakens the others. The solid blue line indicates the breakdown of welfare by component for the 2015 calibration, while the dashed red line shows the



Note: The top left panel plots social welfare as a function of the progressivity parameter, θ_1 , for 2015, under the baseline calibration. Social welfare is measured as the consumption equivalent variation required for agents entering the labor market to be indifferent between the baseline policy and the new one, without accounting for transitions. The vertical lines indicate current and optimal progressivity levels. θ_1^* indicates optimal progressivity according to the model. The right panel shows the social welfare functions for households starting their life in the indicated occupation categories in each year. In this case, welfare is the expected lifetime utility of agents once they have chosen an occupation but before they know a_i (ability) and ϵ_{i1} (starting level of wage risk). Calculations are carried out for the model with no occupational choice.

Figure 12: Optimal progressivity in 2015 and across occupations, without occupational choice. Optimal progressivity increases relative to our baseline experiment.

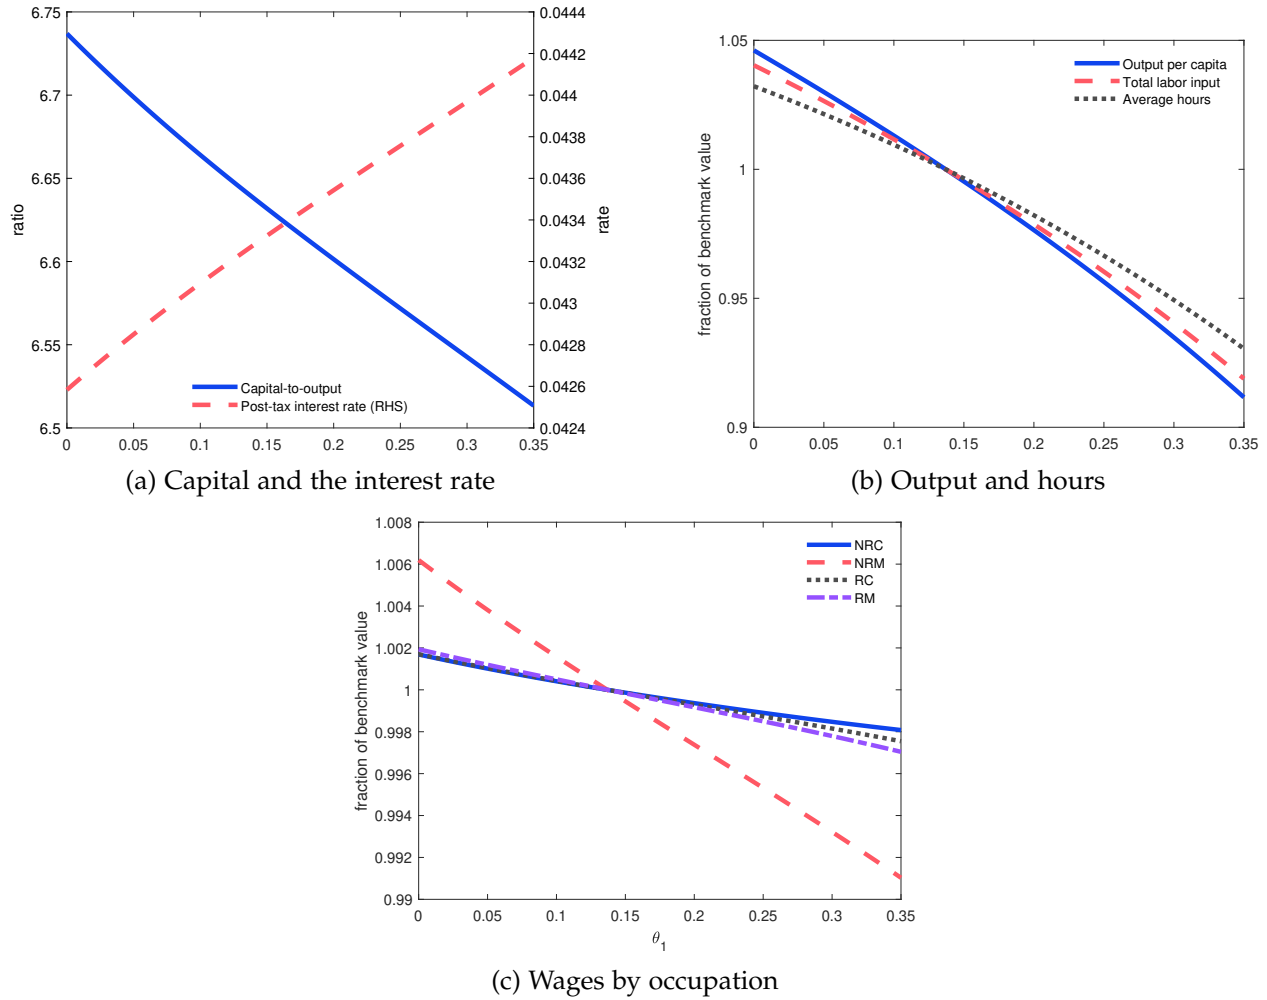
same breakdown but for an economy that did not experience ISTC.

In the latter case, the uncertainty and inequality channels are more sensitive to changes in θ_1 given that the economy has a higher cost of insurance and that the wages of lower-paid occupations react less to an outflow of workers. Additionally, given that no ISTC took place, output per capita is much less sensitive to additional saving by households as a result of lower progressivity.

How would our results change if we ignored the equilibrium effects of occupational choice on wages and other variables? Figure 12 shows how abstracting from occupational choice would have affected our results.

In this case, optimal progressivity in 2015 would be very close to estimated progressivity, though still below it. ISTC still produces a high return rate on capital, which lowers the insurance value of a progressive tax system relative to 1980. The sensitivity of the level channel remains practically unchanged with respect to the economy with occupational choice.

The key difference is on the effects of progressivity on occupational choice and wages.



Note: Total labor input is the sum of labor efficiency units supplied in the economy. Average hours is the percentage of the labor endowment used to work on average. Panels show how prices and quantities in the economy change with respect to progressivity.

Figure 13: Comparative statics of progressivity in 2015 in the absence of occupational choice.

Whereas in the previous exercise employment shares responded to changes in progressivity, in this economy they do not. Hence, a drop in progressivity generates greater capital accumulation and an increase in the intensive margin by NRC workers, but not the extensive margin (Figure 13).

As a result, the net effect of reducing θ_1 depends crucially on the relative strength of capital accumulation versus rising hours in determining wages. Contrary to our previous exercises, lowering progressivity raises both NRC wages and their after taxes earnings, which point to a greater strength of the inequality channel.

This analysis is backed up by the breakdown of the welfare gain by occupation dis-

played on Figure 12b which shows that, in contrast to previous exercises, individuals in the NRC occupation have a preference for no progressivity at all. Meanwhile, agents in the remaining occupations wish for high levels of θ_1 , given that there is no effect positive effect on their hourly wages from reducing progressivity.

In the end, our answer to the question of what is the implication of technological change for optimal taxation does not change in qualitative terms if we abstract from occupational choice, though it does change quantitatively.

8 Conclusion

Since 1980 there has been a steady increase in earnings inequality alongside rapid technological growth in the U.S. economy. To what extent does technological change explain the observed increase in earnings dispersion? How does it affect the optimal progressivity of the tax system? To answer these questions, we develop a life-cycle, overlapping generations model with uninsurable idiosyncratic earnings risk, multiple sources of technological change, a detailed tax system, and occupational choice.

We find that technological transformation can fully account for the increase in earnings dispersion between 1980 and 2015. The main driver is the rising relative wage of non-routine cognitive occupations, which benefit the most from complementarity with capital. In isolation, increasing earnings inequality might strengthen the case for redistributive policies. However, we find that a significant drop in the progressivity of the tax system is optimal. Lower progressivity leads an inflow of workers into higher-paid occupations, raising the wages of those at the bottom of the wage distribution, and incentivises the accumulation of more productive capital, increasing total output. Technological change also raised real return rates on saving, weakening the insurance role of the progressive tax system.

Our work suggests several promising lines for future research. First, while we may find that it is optimal to reduce the progressivity of the labor income tax system, this does not mean that other redistributive policies are not advisable, such as subsidizing access to education and to better paid occupations. How do these policies interact with the tax system, occupational choice and wages?

Second, we do not take into account job displacement due to technological change and non-participation in the workforce. How would this affect our welfare analysis? Is a progressive tax system the right tool to counter these phenomena, or are targeted measures more appropriate?

Finally, our computations do not take into account transitions nor do they look future

technological change. Would a lower progressivity still be optimal in an economy with secular occupation-biased technological growth?

APPENDIX

The Appendix is organized as follows. Section [A](#) indicates micro-data sources and methods. Section [B](#) describes the construction of production factor, price, and output measures. Section [C](#) describes the procedure to estimate the production function. Section [D](#) outlines the procedure for the computation of the equilibrium. Section [E](#) describes welfare evaluation measures and how to decompose them.

A Data sets

A.1 CPS

Imputation. From survey year 1968 to 1975, hours worked in the previous year are not available. We follow [Acemoglu and Autor \(2011\)](#) and impute these by running a regression of hours worked on the previous year on hours worked in the current year, on an indicator variable for whether the individual worked 35+ hours last year or not, on the current labor force status, on an interaction variable between the two previous variables, and on the sector the individual worked in the previous year for the survey years 1976-1978. We then use the estimated equation to assign hours worked in the previous year to the 1968-1975 observations.

Weeks worked last year are not available for 1968-1975 also. We compute mean weeks worked last year by race and gender for the years 1976-78 for each bracket and impute those means for the 1968-1975 period.

Top-coding. To obtain accurate estimates of earnings inequality and wage premia, we have to account for the top-coding in the CPS earnings data. We use the variables *IN-CWAGE*, *INCLONGJ* and *OINCWAGE*, in the taxonomy of [Flood et al. \(2018\)](#). We proceed in two steps: (i) identify top-coded observations; (ii) assuming the underlying distribution is Pareto, we forecast the mean value of top-coded observations by extrapolating a Pareto density fitted to the non-top-coded upper end of the observation distribution. For details on the procedure to approximate the tail of a Pareto distribution see [Heathcote et al. \(2010\)](#).

Top-coding thresholds in the ASEC change across variables and time. Information on top-coding thresholds can be found on the [IPUMS](#) website. Prior to the 1996 survey year, there is little documentation available regarding the thresholds, but the effective top-coding thresholds are provided by IPUMS based on [Larrimore et al. \(2008\)](#). From

1996 onward, the Census Bureau began reporting top-coding thresholds for a set of income variables.

In addition, the Census Bureau has changed its top-coding procedure through time: from 1996 until 2011, the values for top-coded observations were replaced with values based on the individual's characteristics (so-called cell/group means). From 2011 onward, the Census Bureau shifted from an average-replacement value system to a rank proximity swapping procedure.

Ideally, we would like to use a consistent procedure for handling top-coding across time. However, since the Census Bureau started publishing top-coding procedures in 1996, they drastically reduced public use censoring thresholds. [Heathcote et al. \(2010\)](#) found that the Pareto-extrapolation procedure does not perform well in this case. Therefore, we only apply this procedure until survey year 1995. [Heathcote et al. \(2010\)](#) use the extrapolation until survey year 1999, but we find that this produces a large jump in earnings inequality in the late 90's which does not seem plausible.

Bottom-trimming. According to [Flood et al. \(2018\)](#), there is no publicly available information on bottom-coding thresholds of income variables in the ASEC. To deal with this shortcoming, a common practice in the literature is to select a bottom threshold on earnings for inclusion in the sample. We use the procedure of [Heathcote et al. \(2010\)](#): the final sample only includes observations where the hourly wage is above the minimum threshold of one half of the federal minimum wage in each year (end-year federal minimum wage data for farm and non-farm workers is retrieved from FRED).

Variable definitions. All variables are computed as explained in [Acemoglu and Autor \(2011\)](#).

Sample selection. We build two samples, labeled A and B. Table 8 shows the number of records at each stage of the selection process.

The initial sample is a cleaned version of the raw data, which excludes individual records which are either: below the age of 16 in the previous year, not part of the universe, not wage workers, did not work in the previous year, have zero or missing weights, missing age, or have positive earnings but no weeks worked in the previous year, or vice-versa. In 2014, two distinct samples were drawn because of sample redesign. We keep the sample which is consistent with previous surveys.

Sample A excludes all records where the hourly wage is lower than one half of the federal hourly minimum wage. We assume that this sample is representative of the (non-institutionalized) U.S. population. In order to validate the data, we compare a set

Table 8: CPS sample selection (survey years 1968-2017)

	Dropped	Remaining
Initial sample		4,089,617
Wage > 0.5 × min. wage	116,608	3,973,009
Sample A		3,973,009
Age 25-64	861,598	3,111,411
Hours worked per week last year > 6	19,308	3,092,103
Sample B		3,092,103


of sample statistics on wages and hours worked to their aggregate (NIPA) counterpart. This is shown on Figure [A.1](#).

There is an average absolute deviation of 5% between the NIPA (Table 2.1, line 3) and the CPS wage bill. Regarding hours of part and full-time employees, the NIPA series (Tables 6.9B-D, line 2) is lower by 3.3%, on average, and 6.5% after 1986. The BEA uses BLS data to calculate its hours worked series, but the variables are based on the Quarterly Census of Employment and Wages (QCEW) data, rather than on the ASEC variable “usual hours worked per week last year” used in this paper. The total number of full- and part-time employees is much closer to the NIPA series (Table 6.4B-D, line 2), albeit the gap is still 2.7% on average.


Sample B excludes individuals between 25 and 64 years old in the previous year. We consider that 25 years old is a reasonable cutoff age, where individuals’ occupation choice has stabilized. According to the BLS, for 2018 the labor force participation rate drops from 65% to 27%, on average, between the 55-64 and the 65 and older age brackets, which justifies our upper bound for inclusion in the sample. We also exclude records where individuals usually worked less than 6 hours per week in the previous year. This is the sample we use to calculate inequality and wage premia statistics. For comparison, [Heathcote et al. \(2010\)](#) have 2,578,035 individual records in their individual-level database, covering the 1967-2005 survey years. This implies that we have around 63,000 records per year, on average, while [Heathcote et al. \(2010\)](#) have 68,000.

A.2 PSID


Data set structure. The PSID is a panel data set of U.S. individuals and family units. The original 1968 sample was drawn from two independent sub-samples: an over-sample of roughly 2000 poor families selected from the Survey of Economic Opportunities (SEO),



(a) Wages and salaries



(b) Employee full- and part-time hours



(c) Full and part-time employees

Figure A.1: Comparison between aggregate labor variables in the CPS and in the NIPA.

and a nationally representative sample of roughly 3000 families designed by the Survey Research Center (SRC) at the University of Michigan. PSID surveys were annual from 1968 to 1997, and biennial since then.

Since 1968, the PSID has interviewed the individuals from the originally sampled families, which have either remained in the 1968 family unit, or have split off, forming their own. Although some information is collected for each individual in the family unit, the greatest detail is for the so-called husband/reference person and the wife/spouse, when present. In particular, information about wages, occupation, and hours worked are often limited to these two family members, which is the reason why we will focus on these two when analyzing PSID data. See the PSID website for the rules on how the reference person is selected for each family unit.

Because the SRC sample was representative of the U.S. population in 1968, we will restrict our analysis to those families and their split-offs (with a 1968 interview number below 5000). No weights are used for this reason. The main issue with this choice is the inflow of immigrants since 1968. In 1990, the PSID added 2000 Latino households, which covered a major immigration group but missed out on a range of post-1998 immigrants, such as Asians. Because of this short-coming, this sample was dropped in 1995. A new sample of 441 immigrant families, including Asians, was added in 1997 (the so-called “Immigrant” sample).

Variable definitions. To maintain consistency, we use the variable definitions of [Acemoglu and Autor \(2011\)](#), which we used for the CPS data set and which are close to those of [Heathcote et al. \(2010\)](#).

Bottom-trimming. As with the case of the CPS, we eliminate records where the hourly wage is below one half of the end-year federal minimum wage.

Sample selection. As with the CPS, our data cleaning procedure and sample definition procedure is described in this subsection. We build two samples, labeled A and B. [Table 9](#) shows the number of records at each stage of the selection process.

The initial sample is a cleaned version of the raw data on heads and spouses only, and excludes individual records which are: below the age of 16 in the previous year, not wage workers, did not work in the previous year, missing age, or have positive earnings but no weeks worked in the previous year, or vice-versa.

Table 9: PSID sample selection (survey years 1968-2017)

	Dropped	Remaining
Initial sample		260,449
Wage $> 0.5 \times$ min. wage	69,638	190,811
Sample A		190,811
Age 20-64	9,959	180,852
Hours worked last year > 260	5,704	175,148
Sample B		175,148

B Measures

B.1 Labor supply and wages

We follow the procedure of [Krusell et al. \(2000\)](#) to build measures of wages and the labor supply for each of the labor categories (NRC, NRM, RC, RM). The sample used for this purpose is the same as the one used for the regression analysis described on section 3, apart from the fact that we include workers which did not work full-year or full-time. The reason for this is that in the regression analysis we were aiming to identify the wage premia by observing workers in a similar labor market situation. Here, the aim is to construct measures of labor inputs and wages which will be used in the estimation of the production function. We use these bins in order to exclude phenomena such as the increased labor force participation of women from the estimation. Since the labor supply of part-time workers contributes to real GDP, it is necessary to account for those. We do not, however, include self-employed individuals in the analysis. In what follows, the subscript t denotes the year and i denotes an individual observation.

For each worker we record the following variables: hours usually worked per week last year, weeks worked last year, earnings last year, potential experience, race, gender, years of education, occupation category and ASEC weight. Potential experience is divided into 5 five-year groups. Race into white, black and other. There are two sexes. Education is divided into 5 categories: no high school, high school graduate, some college, college graduate, and post college education. Occupation groups are defined as before.

Each worker is assigned to one group defined by the variables described. There are 600 groups, each one denoted by $g \in G$. For each group, we construct a measure of the labor input and labor earnings. The individual labor input is defined as $l_{it} = h_{it}wk_{it}$,

where h_{it} is hours usually worked last year and wk_{it} is weeks worked last year. The individual wage is defined as $w_{it} = y_{it}/l_{it}$. Therefore for each group g we define:

$$l_{gt} = \frac{\sum_{i \in g} l_{it} \mu_{it}}{\mu_{gt}},$$

$$w_{gt} = \frac{\sum_{i \in g} w_{it} \mu_{it}}{\mu_{gt}},$$

where μ_{it} is the individual ASEC weight and $\mu_{gt} = \sum_{i \in g} \mu_{it}$. We aggregate the set G of 600 sets into the occupation categories previously defined $o \in \{\text{NRC}, \text{NRM}, \text{RC}, \text{RM}\}$. From this aggregation we obtain total annual labor input per group, $N_{o,t}$, and its hourly wage, $w_{o,t}$. We assume that the groups within a category are perfect substitutes, and for aggregation we use as weights the group wages of 1980. Thus, for each category o , we have:

$$N_{o,t} = \sum_{g \in s} l_{gt} w_{g80} \mu_{gt},$$

$$w_{o,t} = \frac{\sum_{g \in o} w_{gt} l_{gt} \mu_{gt}}{N_{o,t}},$$

where μ_{it} is the individual ASEC weight and $\mu_{o,t} = \sum_{i \in s} \mu_{it}$. This yields a measure of the total labor input in hours by category ($h_{\text{NRC},t}, h_{\text{NRM},t}, h_{\text{RC},t}, h_{\text{RM},t}$), as well as average hourly wages ($w_{\text{NRC},t}, w_{\text{NRM},t}, w_{\text{RC},t}, w_{\text{RM},t}$).

B.2 Capital, prices and output

Table 10 shows the definitions of main variables compared with those of [Krusell et al. \(2000\)](#).

Table 10: Comparison with [Krusell et al. \(2000\)](#)

Variable	Definition	Definition (KORV)
Output	Business non-farm gross value added	Private domestic product (excluding housing and farm)
Structures	Non-residential structures (private)	Non-residential structures (private)
Equipment	Equipment (private)	Non-military equipment (private)
Equipment price	Equipment price deflator (BEA)	Authors' calculations based on Gordon (1990)

Capital. Our main source for capital data are the BEA's fixed asset accounts and the NIPA. We use only private capital in our measure. Nominal investment for each asset

category is deflated using the investment price index from the BEA.

Equipment prices. To obtain the price of equipment in each year, we aggregate investment price indices from the BEA fixed asset accounts (Table 5.3.4) across equipment types using a Törqvist index. We then divide the resulting average equipment price by the BLS consumer price index for all urban consumers to obtain the relative price of investment.

Depreciation rates. Obtained using the method by [Eden and Gaggl \(2018\)](#). We use BEA data on the net current cost of the stock of capital, $P_{it}NetStock_{it}$, and depreciation at current cost, $P_{it}Dep_{it}$, to compute depreciation rates, which are given by the following formula:

$$\delta_{it} = \frac{P_{it}Dep_{it}}{P_{it}NetStock_{it} + P_{it}Dep_{it}}.$$

We compute average depreciation rates for equipment and non-residential structures, with weights given by the capital stocks at constant prices.

Output. To measure output, we use real gross domestic product in chained 2012 US dollars, retrieved from FRED (FRED code: GDPCA; NIPA code: A191RX).

C Production function estimation method

To estimate the production function, we use the two-step SPML estimator proposed by [Ohanian et al. \(1997\)](#). First, we write the non-linear state space model formally. Next, we briefly describe the methods used to estimate it.

Our non-linear state-space system of equations is of the form:

$$\begin{aligned} \text{Measurement equations :} & & Z_t &= f(X_t, \psi_t, \omega_t; \theta), \\ \text{State equations :} & & \psi_t &= \psi_0 + \psi_1 t + \nu_t. \end{aligned}$$

$f(\cdot)$ contains the labor share equation, the three wage bill equations and the no-arbitrage condition. Z_t is thus a (5×1) vector, which is a function of the variables X_t , the log of the unobservable labor quality indices ψ_t , which is a (4×1) vector, and ν_t and ω_t which are (5×1) and (4×1) vectors, respectively, of i.i.d. normally distributed disturbances. Like [Krusell et al. \(2000\)](#), we assume that A_{t+1} and ψ_{t+1} are known when investment decisions are made.

The model is estimated in two steps: (i) instrument the variables which are potentially endogenous; and (ii) apply the SPML estimator. We assume that the capital stocks, $K_{s,t}$ and $K_{e,t}$, are exogenous at date t . However, we allow for the possibility that date t values of the labor inputs may respond to realization of the technology and labor quality shocks. To instrument these variables, we run a first stage regression of the labor inputs on a constant, current and lagged equipment and non-equipment capital stocks, the lagged relative price of equipment, a trend and the lagged value of the OECD composite leading indicator of business cycles. \tilde{X}_t is the vector of $K_{s,t}$, $K_{e,t}$, the instrumented values of the labor inputs, the depreciation rates and the capital income tax.

The SPML procedure is as follows. Given the distributional assumptions on the error terms, for each t we generate S realizations of the dependent variables, each indexed by i , starting at $t = 1$ in two steps:

$$\begin{aligned} \text{Step 1 :} & \quad \psi_t = \psi_0 + \psi_1 t + \nu_t. \\ \text{Step 2 :} & \quad Z_t^i = f(\tilde{X}_t, \psi_t^i, \omega_t^i, \theta). \end{aligned}$$

In Step 1, we draw a realization of ν_t from its distribution (conditional on our guess of Ω) and use it to construct a date t value for ψ_t . In Step 2, we use our realization of ψ_t , ψ_t^i , together with a draw of ω_t (conditional on our guess of η_ω), to generate a realization of Z_t , Z_t^i . By using this procedure to generate S realization, we can obtain first and second

simulated moments, respectively, of Z_t :

$$m_S(\tilde{X}_t; \theta) = \frac{1}{S} \sum_{i=1}^S Z_t^i,$$

$$V_S(\tilde{X}_t; \theta) = \frac{1}{S-1} \sum_{i=1}^S \left(Z_t^i - m_S(\tilde{X}_t; \theta) \right) \left(Z_t^i - m_S(\tilde{X}_t; \theta) \right)'$$

From this procedure, we will obtain $2T$ moments, which we will use to construct an objective function. Denoting the vector of all actual observations of the dependent variables by Z^T :

$$L_S(Z^T; \theta) = -\frac{1}{2T} \sum_{t=1}^T \left[[Z_t - m_S(\tilde{X}_t; \theta)]' V_S(\tilde{X}_t; \theta)^{-1} [Z_t - m_S(\tilde{X}_t; \theta)] \ln \det(V_S(\tilde{X}_t; \theta)) \right].$$

The SPML estimator, $\tilde{\theta}_{ST}$, is the maximizer of this expression. It is very important that throughout the maximization procedure of the objective function the same set of $(T \times S)$ random realizations of the dependent variables. Otherwise, the likelihood becomes a random function.

D Solution algorithm

To characterize the stationary competitive equilibrium of the model we must find the ratios $\frac{K_s}{N_{NRC}}, \frac{K_e}{N_{NRC}}, \frac{N_{NRM}}{N_{NRC}}, \frac{N_{RC}}{N_{NRC}}$, and $\frac{N_{RM}}{N_{NRC}}$ which clear the capital and labor markets. In addition, we have to fit the tax function, clear the social security budget and find the value of Γ which, given a distribution for the state variable b , uniformly distributes the assets of the dead among the living. G , public consumption of final goods, clears the government budget constraint. The algorithm is as follows:

1. Make a guess on $\frac{K_e}{N_{NRC}}, \frac{N_{NRM}}{N_{NRC}}, \frac{N_{RC}}{N_{NRC}}$, and $\frac{N_{RM}}{N_{NRC}}$.
2. Obtain the value of $\frac{K_s}{N_{NRC}}$ which is consistent with the remaining ratios given the no-arbitrage condition 27 using a bisection method. Compute marginal productivities 17-21 with these guesses.
3. Guess ψ_{ss} , Γ and average earnings.
4. Compute value and policy functions for the retired and active agents by backward induction, given processes for the transitory and permanent shocks. Both shocks are discretized using the Tauchen procedure (Tauchen, 1986), with 4 and 20 states, respectively. We use 20 states for the permanent shock so that we have 5 states for each group supplying a different labor variety. This allows us to calibrate both within-group and between group earnings inequality. The grids for b and n have 24 and 100 points, respectively. In between the grid points, the values of the functions are interpolated using cubic splines.
5. Simulate the model for 120,000 agents, where assets holdings are zero for every agent entering the labor market. Obtain total savings (asset demand), $\int d + \Gamma d\Phi$, and quantities of each labor variety supplied, $N_{NRC}, N_{NRM}, N_{RC}, N_{RM}$.
6. Compute output given the labor supply of households. Asset demand must be allocated between non-equipment and equipment capital. The quantity of structures is obtained by multiplying the initial guess of $\frac{K_s}{N_{NRC}}$ by the quantity of labor supplied by households N_{NRC} . The quantity of equipment, measured in consumption units, is the residual of asset demand. If this residual is negative, we set the quantity of equipment to be 10% of the guess for the non-equipment stock, which allows the algorithm to proceed.
7. Obtain implied values for ψ_{ss} , Γ and average earnings. Compare with guesses made in step 4. If the difference between guesses and implied values is within

a preset tolerance interval, proceed to step 8. If not, update the guesses of each variable and go back to step 4.²³

8. Compute the difference between the ratios implied by the labor supply and asset demand of households with the initial guesses. If these differences are within a preset tolerance level, the solution has been reached with sufficient accuracy. If not, update the guesses and go back to step 2.

E Welfare

A household chooses an occupation after drawing a random vector s which determines the utility of joining each occupation type. Let $\kappa(s)$ denote the idiosyncratic utility of joining a particular occupation given the vector s . Thus, the expected lifetime utility of that household is given by:

$$v(s) = \kappa(s) + \mathbb{E}_0 \left[\sum_{j=1}^J \beta^{j-1} [S_j u(c_j, n_j) + (1 - S_j) D(b_{j+1})] \right]. \quad (\text{A-1})$$

Utilitarian social welfare is defined as:

$$W = \int v(s) d\Phi(s) \quad (\text{A-2})$$

where $\Phi(s)$ is the distribution over the idiosyncratic occupation costs after occupation decisions are made. For a given s , if we scale consumption by $1 + g(s)$ in each period and state of the world, the expected lifetime utility of the resulting allocation is:

$$v(s; g(s)) = \kappa(s) + \mathbb{E}_0 \left[\sum_{j=1}^J \beta^{j-1} [S_j u(c_j(1 + g(s)), n_j) + (1 - S_j) D(b_{j+1})] \right]. \quad (\text{A-3})$$

Isolating $1 + g(s)$ in the right-hand side of the equation:

$$v(s; g(s)) = \kappa(s) + \sum_{j=1}^J \beta^{j-1} S_j \log(1 + g(s)) + \mathbb{E}_0 \left[\sum_{j=1}^J \beta^{j-1} [S_j u(c_j, n_j) + (1 - S_j) D(b_{j+1})] \right]. \quad (\text{A-4})$$

²³Our algorithm uses the homotopy procedure to update all the guesses. That is, if v is the initial guess and v' is the value implied by the simulation, then the updated guess is $v'' = v + a(v' - v)$, where a is a constant chosen by the researcher which controls the size of the update and the rate of convergence of the algorithm.

Integrating over the idiosyncratic state s :

$$W(g) = \log(1 + g) \sum_{j=1}^J \beta^{j-1} S_j + W. \quad (\text{A-5})$$

Let $v(s, B)$ denote the expected lifetime utility of a household with starting state s under an alternative government policy B . Let $W(B)$ denote the expected utility of an unborn individual under that alternative policy. Let $g_U(s)$ denote the consumption equivalent variation necessary to make an individual starting in state s indifferent between benchmark policy A and alternative policy B . These variables must satisfy the following system:

$$v(s; g_U(s)) = v(s, B), \quad (\text{A-6})$$

$$W(g_U) = W(B). \quad (\text{A-7})$$

Given these definitions we can define the consumption equivalent variation for an unborn individual by substituting [A-5](#) into [A-7](#):

$$\log(1 + g_U) \sum_{j=1}^J \beta^{j-1} S_j + W = W(B). \quad (\text{A-8})$$

Solving for g_U :

$$g_U = \exp \left(\frac{W(B) - W}{\sum_{j=1}^J \beta^{j-1} S_j} \right) - 1. \quad (\text{A-9})$$

To breakdown the welfare analysis into *inequality*, *uncertainty* and *level* effects, we decompose g_U using the method of [Flodén \(2001\)](#). We define the certainty-equivalent consumption-leisure bundle for a household starting with state s as:

$$\sum_{j=1}^J u(\bar{c}, \bar{n}) = v(s), \quad (\text{A-10})$$

where \bar{c} and \bar{n} are constant streams of consumption and labor supply. Following [Flodén \(2001\)](#), we set \bar{n} to the average labor supply in the economy. We also remove survival uncertainty on the left-hand side and set the utility of bequests to zero. That leaves only

\bar{c} to be determined. Solving A-10 for \bar{c} :

$$\bar{c} = \exp \left(\frac{v(s)}{\sum_{j=1}^J \beta^{j-1}} + \chi \frac{\bar{n}^{1+\eta}}{1+\eta} \right). \quad (\text{A-11})$$

We can now define the *cost of inequality*:

$$\sum_{j=1}^J \beta^{j-1} u((1 - \rho_{\text{ine}})\bar{C}, \bar{N}) = \int \sum_{j=1}^J u(\bar{c}, \bar{n}) d\Phi(s) = W, \quad (\text{A-12})$$

where \bar{C} and \bar{N} are the average of consumption and labor certainty-equivalents, $\bar{C} = \int \bar{c} d\Phi(s)$, and $\bar{N} = \int \bar{n} d\Phi(s)$. Isolating ρ_{ine} :

$$\rho_{\text{ine}} = 1 - \exp \left(\frac{W}{\sum_{j=1}^J \beta^{j-1}} - u(\bar{C}, \bar{N}) \right). \quad (\text{A-13})$$

The *cost of uncertainty* is defined as:

$$\sum_{j=1}^J \beta^{j-1} u((1 - \rho_{\text{unc}})C, N) = \int \sum_{j=1}^J u(\bar{C}, \bar{N}) d\Phi(s), \quad (\text{A-14})$$

where C and N are average consumption and labor in the economy. Solving for ρ_{unc} :

$$\rho_{\text{unc}} = 1 - \exp (u(\bar{C}, \bar{N}) - u(C, N)). \quad (\text{A-15})$$

Finally, in order to assess the impact of a policy shift in the level of consumption. However, a shift from policy A to B will change equilibrium levels of both consumption and labor. To measure the welfare effects in terms of consumption only, we define *leisure-compensated consumption* denoted by \tilde{C}^B :

$$\sum_{j=1}^J \beta^{j-1} u(\tilde{C}^B, N^A) = \sum_{j=1}^J \beta^{j-1} u(C^B, N^B). \quad (\text{A-16})$$

Solving for \tilde{C}^B :

$$\tilde{C}^B = \exp \left(\log C^B - \chi \frac{N^{B^{1+\eta}} - N^{A^{1+\eta}}}{1+\eta} \right). \quad (\text{A-17})$$

We now have all the ingredients necessary to define the three separate welfare effects of a change from policy A to policy B. Denote g_{lev} as the welfare gain from a change in the levels of consumption and leisure as a result of the policy shift:

$$g_{\text{lev}} = \frac{\tilde{C}^B}{C^A} - 1. \quad (\text{A-18})$$

Denote g_{ine} as the welfare gain from reduced inequality:

$$g_{\text{ine}} = \frac{1 - \rho_{\text{ine}}^B}{1 - \rho_{\text{ine}}^A} - 1. \quad (\text{A-19})$$

Denote g_{unc} as the welfare gain from reduced uncertainty:

$$g_{\text{unc}} = \frac{1 - \rho_{\text{unc}}^B}{1 - \rho_{\text{unc}}^A} - 1. \quad (\text{A-20})$$

[Flodén \(2001\)](#) establishes the following result, which we use in the welfare analysis section to decompose welfare gains into the the three elements:

$$g_U = (1 + g_{\text{lev}})(1 + g_{\text{ine}})(1 + g_{\text{unc}}) - 1. \quad (\text{A-21})$$

References

- Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*. Elsevier.
- Acemoglu, D. and Restrepo, P. (2017). Robots and jobs: Evidence from us labor markets. Working Paper 23285, National Bureau of Economic Research.
- Acemoglu, D. and Restrepo, P. (2018). Artificial Intelligence, Automation and Work. Boston University - Department of Economics - Working Papers Series dp-298, Boston University - Department of Economics.
- Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, 109(3):659–684.
- Aiyagari, S. R. (1995). Optimal capital income taxation with incomplete markets, borrowing constraints, and constant discounting. *Journal of Political Economy*, 103(6):1158–1175.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.
- Benabou, R. (2002). Tax and education policy in a heterogeneous agent economy: What levels of redistribution maximize growth and efficiency? *Econometrica*, 70:481–517.
- Bewley, T. F. (2000). Has the Decline in the Price of Investment Increased Wealth Inequality? Unpublished.
- Brinca, P., Faria-e Castro, M., Ferreira, M. H., and Holter, H. A. (2019). The nonlinear effects of fiscal policy. Working Paper.
- Brinca, P., Holter, H. A., Krusell, P., and Malafry, L. (2016). Fiscal multipliers in the 21st century. *Journal of Monetary Economics*, 77:53–69.
- Cameron, A. and Trivedi, P. (2005). *Microeconometrics*. Cambridge University Press.
- Chamley, C. (1986). Optimal taxation of capital income in general equilibrium with infinite lives. *Econometrica*, 54(3):607–622.
- Cortes, G. M., Jaimovich, N., Nekarda, C. J., and Siu, H. E. (2020). The dynamics of disappearing routine jobs: A flows approach. *Labour Economics*, 65:101823.

- Delaney, K. J. (2017). Droid duties: The robot that takes your job should pay taxes, says bill gates. *Quartz*.
- Eden, M. and Gaggli, P. (2018). On the welfare implications of automation. *Review of Economic Dynamics*, 29:15 – 43.
- Flodén, M. (2001). The effectiveness of government debt and transfers as insurance. *Journal of Monetary Economics*, 48(1):81–108.
- Flood, S., King, M., Rodgers, R., Ruggles, S., and Warren, J. R. (2018). Integrated public use microdata series, current population survey: Version 6.0 [dataset].
- Gordon, R. (1990). *The Measurement of Durable Goods Prices*. National Bureau of Economic Research, Inc.
- Greenwood, J., Hercowitz, Z., and Krusell, P. (1997). Long-Run Implications of Investment-Specific Technological Change. *American Economic Review*, 87(3):342–362.
- Guerreiro, J., Rebelo, S., and Teles, P. (2021). Should Robots Be Taxed? *The Review of Economic Studies*.
- Heathcote, J., Perri, F., and Violante, G. L. (2010). Unequal we stand: An empirical analysis of economic inequality in the united states, 1967–2006. *Review of Economics Dynamics*, 13:15–51.
- Heathcote, J., Storesletten, K., and Violante, G. L. (2019). Optimal Progressivity with Age-Dependent Taxation. CEPR Discussion Papers 13550, C.E.P.R. Discussion Papers.
- Heathcote, J., Storesletten, K., and Violante, G. L. (2020). Presidential address 2019: How should tax progressivity respond to rising income inequality? *Journal of the European Economic Association*, 18(6):2715–2754.
- Holter, H. A., Krueger, D., and Stepanchuk, S. (2014). How Does Tax Progressivity and Household Heterogeneity Affect Laffer Curves? PIER Working Paper Archive 14-015, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania.
- Hugget, M. (1993). The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies. *Journal of Economic Dynamics and Control*, 17:953–969.
- Jordà, Ò., Knoll, K., Kuvshinov, D., Schularick, M., and Taylor, A. M. (2019). The Rate of Return on Everything, 1870–2015*. *The Quarterly Journal of Economics*, 134(3):1225–1298.

- Karabarbounis, L. and Neiman, B. (2014). The global decline of the labor share. *The Quarterly Journal of Economics*, 129(1):61–103.
- Krusell, P., Ohanian, L. E., Ríos-Rull, J. V., and Violante, G. L. (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica*, 68(5):1029–1053.
- Larrimore, J., Burkhauser, R. V., Feng, S., and Zayatz, L. (2008). Consistent cell means for topcoded incomes in the public use march cps (1976-2007). *Journal of Economic and Social Measurement*, 33(2/3):89–128.
- Lucas, R. E. (1990). Supply-side economics: An analytical review. *Oxford Economic Papers, New Series*.
- McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. In Zarembka, P., editor, *Frontiers in Econometrics*, chapter 4, pages 105–142. Academic Press: New York.
- Mendoza, E., Razin, A., and Tesar, L. (1994). Effective tax rates in macroeconomics: Cross-country estimates of tax rates on factor incomes and consumption. *Journal of Monetary Economics*, 34(3):297–323.
- Moll, B., Rachel, L., and Restrepo, P. (2019). Uneven Growth: Automation’s Impact on Income and Wealth Inequality. Boston University - Department of Economics - The Institute for Economic Development Working Papers Series dp-333, Boston University - Department of Economics.
- Ohanian, L. E., Violante, G. L., Krusell, P., and Ríos-Rull, J. V. (1997). Simulation-based estimation of a nonlinear latent factor aggregate production function. In Mariano, R. S., Schuermann, T., and Weeks, M., editors, *Simulation-Based Inference in Econometrics: Methods and Applications*. Cambridge University Press, Cambridge.
- Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economic Letters*, 20:177–181.
- Wu, C. (2021). More unequal income but less progressive taxation. *Journal of Monetary Economics*, 117(C):949–968.