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

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

March 20, 2025

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. We estimate an aggregate production function with capital-occupation complementarity and four occupations that differ with respect to cognitive complexity and routine task intensity. We calibrate our model to resemble the U.S. economy in 1980 and 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. Although technological growth is associated with higher earnings inequality, it leads to a significant drop in optimal tax progressivity. Lower progressivity leads to an inflow of workers into higher-paid occupations. This increases output but also raises the wages of the occupations at the bottom of the wage distribution, dampening the redistributive gains from progressive taxation.

Keywords: Earnings Inequality, Taxation, Technological Change, Automation

JEL Classification: E21, E23, E62, H21, H23, J24, J31, O33, O40

*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 North-American Winter Meeting of the Econometric Society, and the 2019 LubraMacro conference for their helpful comments and suggestions. Pedro Brinca, João Duarte and Hans A. Holter are grateful for financial support from the Portuguese Science and Technology Foundation (FCT), grants number PTDC/EGE-ECO/7620/2020, 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). Pedro Brinca is grateful for his CEECIND/02747/2018 FCT grant, João Duarte for his CEECIND/03227/2018 FCT grant and Hans A. Holter for his CEECIND/01695/2018 FCT grant. João Oliveira is grateful for his FCT grant SFRH/BD/138631/2018. Hans A. Holter thanks the Research Council of Norway, grants: VAM 302661, SKATT 283314, 267428.

†Nova School of Business and Economics.E-mail: pedro.brinca@novasbe.pt

‡Nova School of Business and Economics. E-mail: joao.duarte@novasbe.pt

§University of Delaware, University of Oslo and Nova SBE. E-mail: hans_holter@hotmail.com

¶Bank of England. E-mail: joao.oliveira@bankofengland.co.uk

1 Introduction

Earnings inequality in the U.S. has increased steadily since 1980, see Figure 1 (left panel)[1]. What accounts for the large increase in inequality, and what are the policy implications? There is a heated debate about these questions among academics, policymakers, and the public press. A common view, and perhaps conventional wisdom, is that one should meet increased inequality with higher and more progressive taxes (Delaney, 2017).

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 (ISTC) 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 technological change explain the observed increase in earnings inequality? (ii) how does it affect the optimal progressivity of the tax and transfer system?

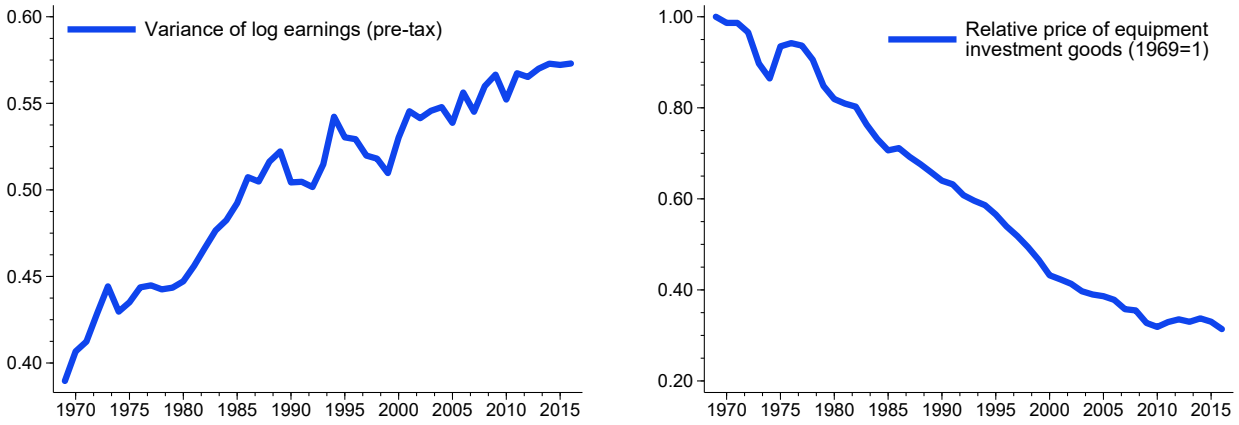
The literature on technological change and the labor market has emphasized task specificity and the degree to which workers' tasks are complementary to capital as crucial determinants of wages. Autor et al. (2003) introduce a framework where occupations differ in terms of the nature of the tasks that 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). To study the evolution of inequality and the implications for optimal tax policy, we adopt this categorization and develop an incomplete markets, heterogeneous agent model with technological change and occupational choice.

Our first contribution is to expand on the seminal paper by Krusell et al. (2000) by specifying and estimating an aggregate production function with labor inputs based on occupation categories rather than the education levels of the workforce. We provide novel estimates for the elasticities of substitution between structures, equipment capital, and these four occupation categories, which have been extensively used in the literature that studies the impact of technological change on labor markets.

Second, and more importantly, we are the first to explain the increase in earnings inequality in the U.S. in a framework with technological growth by estimating a production function with capital-skill complementarity. The previous literature focused on the education skill premium or the labor share of earnings, using representative agent frameworks[2]. We can both explain the changes in wage premia between our four occu-

[1] In 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 the limited availability of other data before 1980.

[2] Krusell et al. (2000) show how capital-skill complementarity can explain the evolution of the college



(a) Inequality

(b) Investment-specific technological change

Note: The variance of log pre-tax earnings is computed from the CPS for employed workers. A description of the data is provided in section A of the Online Appendix. 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.

pation groups and the increase in earnings inequality, measured as the variance of log earnings, between 1980 and 2015[3].

Our third contribution is to investigate the quantitative implications of technological change for optimal tax progressivity in our framework with occupational choice. It is well known that optimal progressive taxation usually comes down to a trade-off between redistribution and insurance on the one hand and efficiency on the other and that introducing margins such as human capital or capital production externalities tilts this trade-off in favor of efficiency. But how strong are these effects? We take a more quantitative approach to this question than the previous literature. We first estimate the aggregate production technology and then study optimal progressivity in our model, which has both human capital (through occupational choice) and a production externality from physical capital (through capital-occupation complementarity), and which furthermore succeeds at explaining the increase in U.S. earnings inequality over time. We show that

wage premium, but they do not study other measures of inequality. Eden and Gaggl (2018) focus on the evolution of the labor share and the routine v.s. non-routine labor share. Finally, Vom Lehn (2020) studies the relative wages of three types of workers as opposed to our four. Using a nested CES production function, where capital equipment is directly substitutable with routine workers only, and a different calibration strategy, he concludes that technological change cannot account for the labor market polarization in wages.

[3]Neither the wage premia nor earnings inequality is targeted by the production function estimation procedure. Explaining them depends on the endogenous savings and labor supply responses of the agents in our model after treating it with the changes in the price of equipment and the levels of technology between 1980 and 2015.

the technological transformation between 1980 and 2015, particularly ISTC, calls for a significant drop in tax progressivity.

Our model is, in some respects, a standard life-cycle model with incomplete markets and idiosyncratic risk. On the household side, it is, however, distinguished by a once and forever choice between our four occupations at the beginning of work-life[4]. Agents make their choice based on an idiosyncratic cost of acquiring the necessary skills and on the expected lifetime utility from consumption and work effort in each profession.

A major departure from previous literature and of crucial importance for the quantitative results of the paper lies in the aggregate production function we use. Our production function has six inputs: four occupations, capital equipment and capital structures. There are three sources of technological growth: ISTC, latent occupation-biased technological change (LAT) and TFP growth. To quantify the labor inputs in each occupation, we apply the cross-walk classification table developed by Cortes et al. (2020) to map tasks into occupation codes. 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. The effect of a fall in the price of equipment investment goods (ISTC) is to spur capital accumulation and create increased demand for workers in occupations with tasks that are more complementary to capital relative to those that are less so. Since there are barriers to mobility between occupations and different entry costs, the rise in labor demand for some occupations creates a wage premium relative to workers in other occupations.

We parameterize the model in two steps. First, we use the firms' first order conditions and a no-arbitrage condition, which restricts the net return of equipment to be the same as of structures, to estimate the production function. We use a simulated pseudo maximum likelihood (SPML) approach, as proposed in Ohanian et al. (1997) and Krusell et al. (2000), that implicitly targets the wage bill ratios of the NRC, NRM and NRM occupations to the RM occupation from 1968 to 2015. Our estimates deliver not only a good fit to the targeted moments but also a good fit to two *non-targeted* data moments: the wage premia of each of the three occupations, NRC, NRM and NRM, relative to the RM occupation[5] and the growth rate of total factor productivity from 1968 to 2015. Second, we insert the estimated production function into our incomplete markets model, and we calibrate the remaining parameters to resemble the U.S. economy in 1980.

[4]Some workers do of course retrain; however, Cortes et al. (2020) provide evidence of the fall in routine employment in the U.S. being primarily caused by declining inflow rates among younger workers.

[5]This is using quantities of labor and capital from the data. It will also turn out to work in our model with endogenous savings and labor supply decisions.

Inserting the growth of ISTC, LAT and TFP, we find that technological change (in particular ISTC) can fully account for the changes in wage premia and increase in earnings inequality between 1980 and 2015[6]. Neither the wage premia nor earnings inequality is targeted by the production function estimation procedure, and explaining them is dependent on the endogenous savings and labor supply responses of the agents in our model after treating it with the changes in the price of equipment (ISTC) and the levels of LAT and TFP between 1980 and 2015. The main driver of the increase in earnings inequality 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. ISTC alone accounts for about 2/3 of the increase in earnings dispersion, and latent occupation-biased technological change accounts for the remaining 1/3.

Our optimal tax experiment is to maximize the expected steady-state welfare of an unborn individual with respect to the progressivity and level of the labor income tax code, taking government expenditure and other taxes as exogenously given[7]. We then study the interaction between optimal tax progressivity and our three sources of technological growth, and we use the framework of Flodén (2001) (see also Benabou (2002) for a similar approach) to decompose the welfare effects of progressive taxation into the contributions resulting from its impact on efficiency, redistribution and insurance.

We apply a non-linear tax function as in Benabou (2002) and Heathcote et al. (2017), $y_a = 1 - \theta_0 y^{-\theta_1}$, where y_a denotes after-tax income and θ_0 and θ_1 define the level and progressivity of the tax system, respectively. For 1980, we find the optimal value of our measure of tax progressivity, θ_1, to be 0.15 (close to the estimated benchmark value of 0.19), whereas, in 2015, a value of 0.05 is optimal[8]. To give an interpretation in terms of actual tax rates: The average tax rate for an individual with Average Earnings (AE) is 15% both with $\theta_1 = 0.15$ and $\theta_1 = 0.05$. The average tax rates for two individuals making 0.5AE and 2AE are, however, 5.7% and 23.4% with $\theta_1 = 0.15$ and 12.0% and 17.9% with $\theta_1 = 0.05$. The main mechanisms driving this result are the high productivity of NRC professions in 2015, the positive effect of shifting workers to NRC occupations

[6]This finding is consistent with Barro (2000) who finds that across rich counties, inequality and economic growth are correlated.

[7]This is the classic tax experiment in the literature on incomplete market models with heterogeneous agents. The recent literature also studies transitions, but given the complexity of our model relative to this literature (we do, for example, have to solve for five prices in equilibrium), we are restricted to studying steady states for now. Recent contributions to the optimal taxation literature in macroeconomics with transitions such as Boar and Midrigan (2022), Ferriere et al. (2022) work with the standard Aiyagari (1994) model.

[8]Indeed, there is evidence of some reduction in tax progressivity in the U.S. since 1980. Wu (2021) finds that this measure of progressivity has fallen from 0.19 to 0.14 between 1980 and 2015.

on the wages of lower-paid occupations, and the higher returns to wealth with the 2015-technology[9]. Reducing tax progressivity shifts workers towards higher-paying occupations[10], which raises output as well as the wages in lower-paying occupations, but also reduces the benefits of redistribution and insurance from the tax system. This tradeoff is, however, tilted towards flatter taxes with the technological transformation between 1980 and 2015.

Among our three sources of technological growth, ISTC is solely responsible for the drop in optimal tax progressivity (LAT and TFP growth pulls in the other direction). From the perspective of the social planner, all three welfare impacts of progressive taxation (efficiency, redistribution and insurance) are tilted towards lower optimal progressivity with higher ISTC. First, the efficiency channel is stronger because there is more capital and stronger complementarity with high-earning professions. The benefit from lowering the marginal tax rates on high earners and getting people to select NRC professions is thus higher. Second, although there is more earnings inequality in 2015, which creates additional incentives for redistribution, more agents moving from low-earning to high-earning occupations increases the wage rates of low earners and decreases the wage rates of high earners. The positive effects that people moving to high-earning occupations have on the wages of low-earning occupations dampens the redistributional loss from flatter taxes. Finally, ISTC is responsible for the increased returns on capital in 2015, which dampens the insurance motive. A higher return on capital makes it easier to self-insure and weakens the insurance role of a progressive tax system.

The rest of the paper is organized as follows. Section 2 contains a brief survey of the related literature. In Section 3, we describe the model. In Section 4, we estimate our aggregate production function. Section 5 is devoted to calibrating our model. In Section 6, we present our quantitative results on inequality and optimal taxation. We conclude in Section 7.

2 Relation to the Literature

This paper relates to two main strands of literature. First, the literature investigating the impact of technological change on wages and inequality and, second, the literature on optimal Ramsey taxation in incomplete markets models with heterogeneous agents.

[9]See Jordà et al. (2019) for evidence of higher return rates on wealth in the U.S. Moll et al. (2019) also argue that technological growth raises the return on wealth.

[10]Without occupational choice there is only a slight drop in optimal progressivity between 1980 and 2015, see Section G of the Online Appendix

Our work builds on the classic paper by Krusell et al. (2000). We expand their framework by specifying and estimating an aggregate production function with labor inputs based on occupations rather than the education levels of the workforce. Krusell et al. (2000) document the impact of skill-biased technological change and capital-skill complementarity on the skill premium (i.e., the college premium) and can explain its evolution over time using this mechanism. Their approach is, however, a purely production-side approach with two types of labor (high-skilled and low-skilled). They do not model households' endogenous savings and labor supply decisions, and they do not study other inequality measures, such as the variance of earnings. Using our framework with four types of labor but also rich agent heterogeneity in the forms of income risk, age, savings and permanent ability, we can both explain the changes in skill premia between our four occupation groups as well as the increase in earnings inequality in the U.S., measured as the variance of log earnings between 1980 and 2015. This result depends on our estimation of the production function and the endogenous savings and labor supply decisions of the agents in our model in response to the changes in the price of equipment (ISTC) and the levels of LAT and TFP.

Instead of dividing the population by education level, Autor et al. (2003) argues that the most empirically relevant interaction between technology and worker productivity comes from the types of tasks a worker performs (although these are correlated with education). They study the effect of computerization on changes in employment by occupation categories and posit that some occupations have a prevalence of tasks that can easily be automated and solved by machines (routine tasks). In contrast, others involve complex problem-solving and interactions (so-called non-routine tasks) which are very costly or impossible to automate. The other key distinction of tasks is whether they are cognitive or manual. We adopt the occupation taxonomy of Autor et al. (2003) and use the cross-walk classification table developed by Cortes et al. (2020) to map tasks into occupation codes to calculate equilibrium quantities of labor input by occupation category[11].

There is a growing literature classifying labor inputs by tasks and studying the interaction with automation technologies. 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. Our work instead uses the four task dimensions postulated by Autor et al. (2003). Also, it allows for labor-augmenting technological change at the occu-

[11]See Online Appendix A for additional details on data treatment. We use these data to construct 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 Online Appendix B.

pation level, which will be important for our findings below showing that workers at the bottom of the wage distribution have enjoyed wage growth relative to the center of the distribution as a result of technological change. Vom Lehn (2020) maps tasks into three labor types and proposes an aggregate function that is closer, but still quite different[12], from our aggregate production function specification. In contrast to our results, he finds that his calibrated model cannot reproduce the job market polarization in wages. The difference between his findings and ours possibly stems from the different production function specifications, the calibration procedures or the different classifications of labor inputs. Other papers using a task-based framework to study the impact of technological growth on inequality include Acemoglu and Autor (2011), Acemoglu and Restrepo (2018), Moll et al. (2019), Kaplan and Zoch (2020). We do not follow some of these studies in modeling tasks explicitly. We thus forego a more detailed characterization of the production process in favor of the ability to measure the inputs in production more accurately, enabling the estimation of the production technology in Section 4 below.

This paper is also related to the literature on optimal progressive Ramsey taxation in incomplete markets models with heterogeneous agents. Due to the complexity of our model we focus on maximizing steady state welfare, as in Conesa and Krueger (2006), Conesa et al. (2009), Peterman (2016), Heathcote et al. (2017), Heathcote et al. (2020), Wu (2021). In the same tradition, there is also a recent sizeable literature considering transitions after once and forever tax changes, see e.g. Bakis et al. (2015), Kindermann and Krueger (2022), Boar and Midrigan (2022), Ferriere et al. (2022) and a much smaller literature studying optimal dynamic taxation during a transition, see e.g. Acikgoz et al. (2022), Dyrda and Pedroni (2021). All these papers with transitions do, however, have in common that they work with the classical Aiyagari (1994) model. Our contribution is to quantify the impact of technological change and human capital (through occupational choice) on optimal tax progressivity. These are two factors of crucial importance to inequality as well as the trade-offs between efficiency, redistribution and insurance that one must consider when designing optimal tax systems.

Some recent studies have raised the question of how the tax system should respond to increasing inequality caused by various sources. Closest to ours are Wu (2021) and Heathcote et al. (2020). Wu (2021) considers an ageing population, shrinking gender wage gap, increased idiosyncratic risk, and increased skill premium (modeled with a parameter governing the returns to human capital investment). In total, these changes lead to a slight drop in optimal tax progressivity. The effect of an increase in the skill premium

[12]In his production function, abstract and manual labor inputs are substitutes or complement to a bundle composed of routine labor input and capital equipment. In contrast, in our framework, NRC, NRM and RC all have a constant elasticity of substitution with capital equipment directly.

(he captures it with a parameter governing the returns to human capital investments) on optimal progressivity is, however, almost neutral. Heathcote et al. (2020) study the impact of technological change on optimal progressivity in an incomplete markets model with skill choice. They also find that skill-biased occupational choice is almost neutral concerning optimal tax progressivity. However, their focus is on college education and skill-biased technological change, and there is no role for capital in production. Our paper takes an occupation-based approach and focuses on the role of capital-occupation complementarity. In contrast to these two studies, we find a striking drop in optimal tax progressivity due to ISTC.

Related to our work is also Ales et al. (2015) who study Mirrlesian taxation in a talent assignment model in a static model without capital but with technical change. They find that technical change should lead to a slightly more progressive tax system. Scheuer and Werning (2015) study the impact of superstars on optimal Mirrlesian taxation. They find the impact of superstars on the optimal tax system to be neutral. Guerreiro et al. (2022), study optimal capital taxation in a less quantitative model than ours but with the possibility of automation of tasks and endogenous choice between two skills/occupations[13]. In their model, which has fewer state variables than ours, it is possible to study transitional dynamics. They find that one should tax robots in the short run but not in the long run, even if it leads to higher inequality. Our contribution is distinct from theirs in that we take the model to data and look at the implications of the changes in the price of equipment and technology that has taken place over time for optimal taxation. We also broaden the analysis to include the cognitive/manual dimensions of tasks, and we include idiosyncratic income risk (adding an insurance motive to the optimal taxation problem).

Finally, our paper which has human capital investments modeled as occupational choice relates to the literature studying the impact of human capital investments on inequality and the interaction with government policy, such as Huggett et al. (2011), Guvenen et al. (2014), Holter (2015), Herrington (2015), Badel and Huggett (2017). In our case, like in most of the literature, we find that reduced tax progressivity leads to higher inequality; however, this is not necessarily negative as long as most households (including the bottom earners) become richer.

[13]Like us they assume that older generations cannot change occupations. This 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.

3 A Model of Labor Market Inequality and Technological Change

Our model is a life-cycle version of the Bewley-Aiyagari-Hugget model:[14] 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 consumption and leisure.

Before 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 throughout its lifetime.[15] After labor market entry, households face a stream of idiosyncratic wage shocks 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: the consumption goods, structure capital goods, and equipment capital goods sectors. 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 consumer goods sector, which is the formulation that Krusell et al. (2000) adopts to incorporate investment-specific technological change.

The centerpiece of the model is the production function for the intermediate input, which uses a combination of the different occupation and capital types to produce final goods. We build on the production function introduced by Greenwood et al. (1997) and extend it to encompass 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 the capital and labor demand and 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 of these variables into earnings inequality measures.

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

[14] See Bewley (2000), Aiyagari (1994), and Hugget (1993).

[15] Cortes et al. (2020) provide evidence of the main driver of the decline in routine employment being a reduction in inflow rates rather than an increase in outflow rates. This is consistent with our assumption of an inability to change occupation type in the middle of working life, despite changing wage premia in other occupation types.

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

Below, we describe the household problem, the production side of the economy, and the definition of equilibrium in more detail.

3.1 Demographics

We assume the economy is populated by a set of $J = 81$ overlapping generations, as in Brinca et al. (2016). A period in the model corresponds to one year, and households begin life at age 20. Thus, j, the household's age, varies between 0 (for age 20 households) and 80 (for age 100 households).

Before joining the labor market, agents must make an irreversible and mutually exclusive occupation choice, deciding which labor market will determine their wages throughout 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 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).[16] 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.[17] Other than occupation, households differ in the value of their persistent,

[16] Concretely, 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.

[17] To find V_o for each occupation, we calibrate and solve a version of the model where occupations

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

After retiring at age 65 (model age 45), 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, so that in any given period, using a law of large numbers, the mass of retired agents of model-age $j \geq 45$ is equal to $S_j = \prod_{t=45}^{t=j} s_{t-1}$.

Dying households leave bequests which are redistributed evenly 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. (2021).

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 set to ensure that the Social Security system breaks even in equilibrium.

3.2 Preferences

The momentary utility function, $u(c_{ij}, n_{ij})$, depends on consumption, c_{ij}, and labor supply, $n_{ij} \in (0, 1]$, and is given by:[18]

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

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

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

are randomly assigned in such a way that we match the employment weights of each occupation type in 1980. The 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.

[18] We assume that disutility of work depends only on working time, 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 the occupational choice is given by:

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

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

3.3 Labor Income

Labor productivity depends on three elements that determine the efficiency units of labor 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:

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

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

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

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, to match within-occupation earnings dispersion. Online Appendix D describes the implementation of this procedure in the numerical algorithm.

3.4 Technology

In this framework, there exist three competitive final goods: consumption goods, structure investment goods, and equipment investment goods. Each 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 households and the government.

The transformation technology is:

$$C_t + G_t = Z_{c,t}, \tag{8}$$

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:

$$p_{c,t} = 1 = p_{z,t}. \tag{9}$$

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

$$X_{s,t} = Z_{s,t}, \tag{10}$$

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

$$X_{e,t} = \frac{Z_{e,t}}{\xi_t}, \tag{11}$$

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

$$p_{e,t} = \xi_t p_{z,t} = \xi_t, \tag{12}$$

where $\xi_t = p_{e,t}/p_{c,t}$ 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, $y_t = F(K_{s,t}, K_{e,t}, N_{NRC,t}, N_{NRM,t}, N_{RCt}, N_{RM,t})$, 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: $Y_t = C_t + G_t + X_{s,t} + \xi_t X_{e,t}$, factor prices, and the price of the intermediate good $p_{z,t}$ are taken as given. The firm chooses capital and labor inputs for each period to maximize profits:

$$\Pi_{z,t} = p_{z,t} y_t - r_{s,t} K_{s,t} - r_{e,t} K_{et} - \sum_{o \in O} w_{ot} N_{ot}, \tag{13}$$

subject to:

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

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 $F(.) = Y_t = C_t + G_t + X_{s,t} + \xi_t X_{e,t}$. We assume that the production function of intermediate goods is Cobb-Douglas over structure capital and CES over the remaining inputs:

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

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

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

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[19]. 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. Firm maximization implies that marginal products equal factor prices[20].

[19] 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 extends this framework with a finer breakdown of 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.

[20] The first-order conditions can be found in section F of the Online Appendix.

Capital laws of motion are given by:

$$K_{s,t+1} = (1 - \delta_s)K_{s,t} + X_{s,t}, \tag{16}$$

$$K_{e,t+1} = (1 - \delta_e)K_{e,t} + X_{e,t}, \tag{17}$$

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

3.5 Government

The social security system is managed by the government and runs a balanced budget. The 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), Heathcote et al. (2017) and Holter et al. (2019):

$$y_a = 1 - \theta_0 y^{-\theta_1}, \tag{18}$$

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.[21]

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:

$$T_t = G_t, \tag{19}$$

$$\Psi_t \left(\sum_{j \geq 45} \Omega_j \right) = R_t^{ss}. \tag{20}$$

3.6 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., $\xi_{t+1} = \xi_t = \xi$, so we drop the time index on return rates for this exposition. Thus, the return rates

[21] See the Online Appendix of Holter et al. (2019) for a detailed discussion of the properties of this tax function.

must satisfy:

$$\frac{1}{\xi}\left[\xi + (r_e - \xi\delta_e)(1-\tau_k)\right] = 1 + (r_s - \delta_s)(1-\tau_k), \tag{21}$$

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:

$$b_{ij} \equiv \xi k_{e,ij} + k_{s,ij}, \tag{22}$$

3.7 Household Problem

In any given period a household is defined by its 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 of maximizing expected utility. The household budget constraint is given by:

$$\begin{aligned} c_{ij}(1+\tau_c) + \xi k_{e,ij+1} + k_{s,ij+1} = \left[\xi + (r_e - \xi\delta_e)(1-\tau_k)\right] k_{e,ij} \\ [1 + (r_s - \delta_s)(1-\tau_k)]k_{s,ij} + q\Gamma + Y^N, \end{aligned} \tag{23}$$

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 21, in equilibrium we can rewrite the budget constraint as:

$$c_{ij}(1+\tau_c) + b_{ij+1} = (b_{ij} + \Gamma)[1 + r(1-\tau_k)] + Y^N. \tag{24}$$

The household problem can be formulated recursively as:

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

s.t.:

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

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

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

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

retired household's problem can be written as:

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

s.t.:

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

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

3.8 Stationary Recursive Competitive Equilibrium

Letting $\Phi(j, b, o, a, u)$ be the measure of agents with corresponding characteristics (j, b, o, a, u), we define a stationary recursive competitive equilibrium as follows[22]:

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:

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

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

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

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

3. The prices of the production factors equal their marginal products (Equations A-24-A-28 hold).

4. The government budget balances:

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

5. The social security system balances:

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

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

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

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

4 Estimating the Production Function

The production function in Equation 15 that transforms our four labor varieties and two types of capital into output goods is crucial to our quantitative results. In this section, we describe the stochastic specification of the production function model, the equations to be estimated, and the results. The estimation strategy follows Krusell et al. (2000). When we calibrate our model, we will treat the parameter estimates from this section as exogenously given. When we study the impact of changes in technology over time on inequality, we will insert our results from this section in the model[23].

4.1 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:

$$\psi_t \equiv \log(\varrho_t), \quad \psi_t = \psi_0 + \psi_1 t + \nu_t, \tag{25}$$

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: $\nu_t \sim N(0, \Omega)$. The i.i.d. assumption simplifies the identification of the factor-neutral technological change, A_t, which is described below.

4.2 Equation Specification

We use a system with two sets of equations obtained from the first order conditions of agents to estimate the model: (i) the wage bills relative to the routine manual labor

[23] The data used in the estimation is described in Online Appendix B.

variety; and (ii) a no-arbitrage condition between investing in equipment and structure capital. These are defined as follows:

$$\frac{w_{o,t}h_{o,t}}{w_{\text{RM},t}h_{\text{RM},t}} = wbr_{o,t}(\psi_t, X_t; \theta), \qquad o \in O = \{\text{NRC}, \text{NRM}, \text{RC}\}, \tag{26}$$

and

$$1 + [F_{K_s}(\psi_{t+1}, X_{t+1}; \theta) - \delta_{s,t+1}] = E_t\left(\frac{\xi_{t+1}}{\xi_t}\right)(1 - \delta_{et+1}) + \frac{F_{K_e}(\psi_{t+1}, X_{t+1}; \theta)}{\xi_t} \tag{27}$$

where equation 27[24] is obtained from equation 21, assuming that $\xi_t \neq \xi_{t+1}$, and where we substituted the return rates by factor marginal productivities.

Depreciation rates are indexed by t since they change over the time[25]. The relative wage bills in the model $wbr_{o,t}$ 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 27, and that the tax treatment is identical between equipment and structure capital returns. Finally, we substitute the first term on the right hand side of equation 27 with $E_t\left(\xi_{t+1}/\xi_t\right)\left(1 - \delta_{et}[1 - \tau_{k,t}]\right) + \omega_t$, where ω_t is the i.i.d. forecast error and $\omega_t \sim N\left(0, \eta_\omega^2\right)$. This set of assumptions imply that $A_t = Y_t/G(.)$ from equation 15.

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[26].

The parameter vector θ has dimension 36. Our sample contains 49 observations for each equation. We reduce the number of parameters estimated 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 is identical for all labor types. Thus, $S = \eta_\nu^2 I_4$, where η_ν^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

[24]Note that this no-arbitrage equation applies on capital returns net of depreciation. Hence, in equilibrium, we are allowing for different capital gross returns across the two types of capital because they have different depreciation rates.

[25]See Online Appendix B

[26]See details in Online Appendix C.

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 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_ν^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, $K_{e,0}$.

4.3 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 the 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, they estimate an elasticity of substitution of 1.4 for non-routine labor. 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. For the processes of occupation-specific technology, we estimate that only the non-routine cognitive occupations have experienced positive growth. At the same time, routine manual labor has suffered the largest decline[27].

Figure 2 shows model fit to targeted moments over time. Figure 2a displays aggregate *ex post* return rates of equipment and structures implied by our model, which are zero 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 2b plots wage bill ratios implied by the model, as specified by the set of equations (26), and the data. Model predictions closely track the data. The NRC wage bill

[27] Vom Lehn (2020) also estimates elasticities of substitution between different task bundles that are not directly comparable to ours. In his production function, abstract and manual labor inputs are substitutes or complement to a bundle composed of routine labor input and capital equipment. In contrast, in our framework, NRC, NRM and RC all have a constant elasticity of substitution with capital equipment directly. In the case of homogeneous workers, he calibrates the elasticity of substitution between routine labor input and capital equipment to 1.3, between manual and a bundle of routine labor input and equipment to 1.49, and between abstract labor input and a bundle of routine labor input and equipment to 0.31.

Table 1: Parameter 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 NRM	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

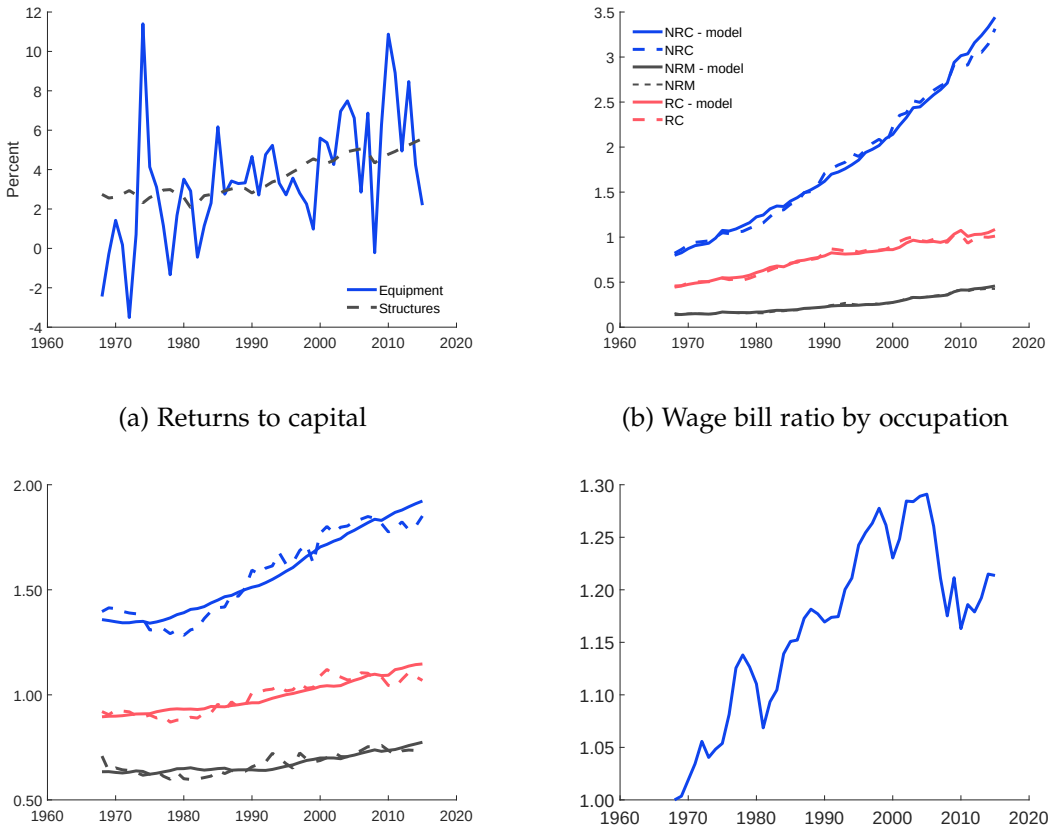
Note: The table shows the parameter estimates for the production function and the labor efficiency indices. "EOS" stands for elasticity of substitution. The ϕ are the shares of each occupation inside each labor-equipment composite. The φ are the shares of each labor-equipment composite. The ψ_0 indicate the intercept of the linear labor efficiency indices, and ψ_1 the slope. $K_{e,0}$ is the starting level of equipment capital in millions of dollar.

shot up from near 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 2c 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 are 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. The key force driving earnings dispersion is the change in wage premia across groups.

Finally, Figure 2d 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).

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



(a) Returns to capital

(b) Wage bill ratio by occupation

(c) Wage premia (non-targeted)

(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 to estimate the model. In Figure 2d, total factor productivity is normalized to 1 in 1968. Construction of the measures is described in Online Appendix B.

Figure 2: Empirical Model Fit to Targeted and Non-Targeted Moments.

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.

5 Calibration

This section describes the calibration of the baseline model to resemble the U.S. economy in 1980. Many parameters can be set externally (i.e., we estimate them directly from the data or take them from the previous literature and insert them in the model). This includes the production function that we estimated in Section 4 but also, for example,

the tax function and the age profile of earnings. Table 2 lists the externally calibrated parameter values and data sources. The seven parameters in Table 4 are estimated by the simulated method of moments (SMM) approach, where we change the parameters to minimize the distance between the model and data moments.

5.1 Externally Calibrated Parameters

Below we discuss the external calibration of parameters that were not estimated in 4.

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:[28]

$$\ln(w_{it}) = a_i + \gamma_1 j + \gamma_2 j^2 + \gamma_3 j^3 + \varepsilon_{it}. \tag{28}$$

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 3. 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 Online 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, $\tau_c = 0.05$, $\tau_k = 0.47$.

[28] PSID data is described in section A of the Online Appendix.

Table 2: Externally Calibrated Parameters

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 4
Structures depreciation rate	δ_s	0.026	Section 4
Share structures	α	0.040	Section 4
Share NRC	ϕ_1	0.378	Section 4
Share NRM	ϕ_2	0.086	Section 4
Share RC	ϕ_3	0.279	Section 4
Share composite NRC	φ_1	0.160	Section 4
Share composite NRM	φ_2	0.045	Section 4
Share composite RC	φ_3	0.023	Section 4
EOS NRC	ρ_1	0.497	Section 4
EOS NRM	ρ_2	2.055	Section 4
EOS RC	ρ_3	5.029	Section 4
EOS RM	σ	5.564	Section 4
Latent efficiency NRC	ϱ_1	2.734	Section 4
Latent efficiency NRM	ϱ_2	4.955	Section 4
Latent efficiency RC	ϱ_3	34.662	Section 4
Latent efficiency RM	ϱ_4	0.378	Section 4
Total factor productivity	A	16.728	Section 4
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

5.2 Endogenously Calibrated Parameters

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

$$L(\tilde{\theta}) = ||M_m - M_d||, \tag{29}$$

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 (29).

Table 3: Fit of Model to Data (SMM)

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
$\overline{n}$	Fraction of hours worked	BEA	1/3	1/3
$\text{Var}\ln(e_{\text{NRC}})$	Variance of log earnings (NRC)	CPS	0.408	0.409
$\text{Var}\ln(e_{\text{NRM}})$	Variance of log earnings (NRM)	CPS	0.410	0.406
$\text{Var}\ln(e_{\text{RC}})$	Variance of log earnings (RC)	CPS	0.409	0.410
$\text{Var}\ln(e_{\text{RM}})$	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
$\sigma_{a,\text{NRC}}$	0.519	Variance of ability NRC
$\sigma_{a,\text{NRM}}$	0.515	Variance of ability NRM
$\sigma_{a,\text{RC}}$	0.517	Variance of ability RC
$\sigma_{a,\text{RM}}$	0.385	Variance of ability RM

We use the ratio between the 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[29]. The capital stock is obtained from the estimation of the empirical model of section 4. Disutility from work targets average 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. Table 4 presents the

[29]This quantity is not directly comparable to the usual K/Y in the macro literature because we include non-residential capital in the form of non-residential structures and equipment capital only, as proposed by Krusell et al. (2000).

parameters calibrated internally through SMM estimation, and Table 3 displays the fit of the model moments to the data moments.

6 Quantitative Results

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

6.1 The Sources of Growing Earnings Inequality

The main experiment conducted in this section is to change the externally estimated levels of technology and parameters governing the tax system from their values in 1980 to their 2015 values[30]. In addition, we recalibrate the distribution of occupation-specific utilities to match the occupation shares in 2015. We then decompose the variation in our earnings inequality measure (the variance of log earnings) between the two steady states to identify the role of investment-specific technological change (ISTC), latent occupation-biased technological change (LAT), and TFP growth.

We also compare the magnitude of these technological sources of variation in earnings dispersion to others, such as the observed changes in the progressivity of the tax system and the fall of the capital income tax. 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 (γ_1, γ_2, γ_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. The parameters changed from 1980 to 2015 are listed in Table 5.

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 4. Likewise, TFP is set to equal the estimated level in

[30]There many other factors that potentially changed between 1980 and 2015, and that could cause either higher or lower inequality (for example, ageing population, shrinking gender wage gap, and changes in idiosyncratic risk, see Wu, 2021) but we focus on these two channels.

Table 5: Parameter Changes 1980-2015

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

2015. The location parameters of the idiosyncratic cost distributions are set to 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 estimates of Wu (2021). The Social Security tax rates are those described in Brinca et al. (2016) for the U.S. economy. 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 some untargeted data moments in 1980 and 2015. In the first section of the table, we compare relative input quantities from the theoretical model to those obtained from estimating the empirical model of the production function in section 4. The relative input quantities are fairly close to our estimates, except for the growth of equipment capital between 1980 and 2015, which the theoretical model substantially underestimates relative to the empirical model. In other words, the model cannot generate a sufficiently large rise in savings compared to our empirical estimates.[31]

The second section shows the wage changes by occupation between both steady states. The model slightly overestimates wage growth for all occupations, except for NRM. However, as can be seen in the third section of the table, wage premia are very close to the data in both years, which is key in terms of accounting for the change in earnings dispersion. The bottom line shows the variance of log earnings which is the

[31]This is likely because the U.S. is turning into a large open economy. This means that the stock of capital has grown via an increase not only in domestic savings but also in 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). See also Chakraborty et al. (2017) for the growth in cross-border lending to U.S. firms over the period.

Table 6: Model Fit

Variable	1980 Model	1980 Data	2015 Model	2015 Data
Relative input quantities				
K_e / N_{NRC}	6.27	7.80	16.70	39.14
N_{NRM} / N_{NRC}	0.66	0.55	0.49	0.43
N_{RC} / N_{NRC}	10.61	9.10	5.74	5.66
N_{RM} / N_{NRC}	0.16	0.16	0.05	0.05
Wage growth				
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
Wage premia				
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
Variance of log earnings	0.43	0.45	0.57	0.57

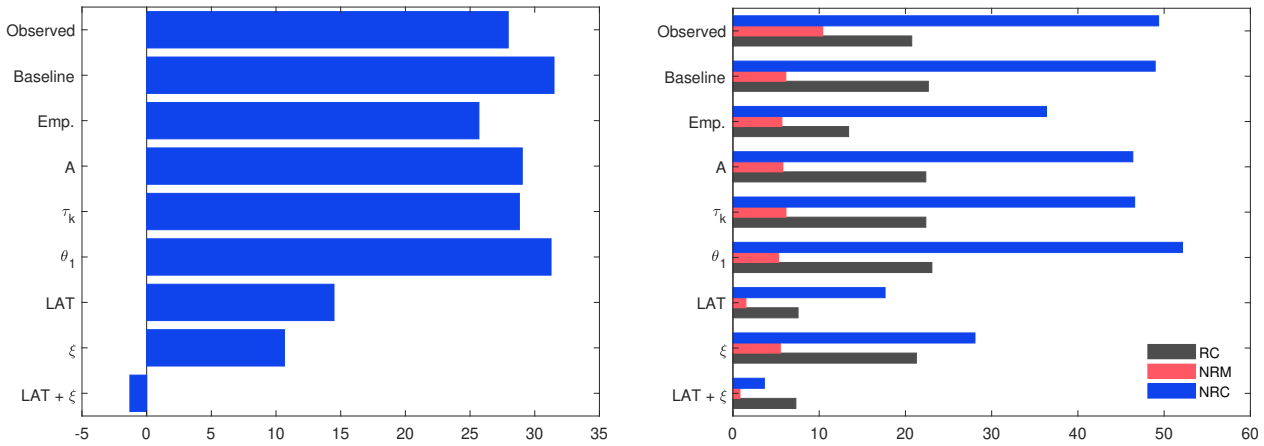
Note: The first section of the table displays the relative input quantities. The second section displays wage per efficiency unit by occupation (model definition) and wages per hour at constant 1968 prices by occupation (data definition). All the prices are normalized to 1 in 1980. The third section shows the wage premia calculated as the ratio between the marginal productivities of labor in 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.

centerpiece of our analysis. The total variance of log earnings grows 27% from 1980 to 2015 in the data and 33% in the model, implying that the model slightly overshoots the growth in earnings inequality.

To understand the drivers of the change in aggregate earnings inequality, we decompose the model predicted 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 3 illustrates these exercises by displaying the response of labor income dispersion measures to the set of parameter shifts presented in 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 due to the baseline parameter shifts shown in Table 5.

Figure 3a, in the top left panel, shows how the model fares in generating a shift



(a) Change in variance of log pre-tax earnings (%)

(b) Change in wage premia (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 A of the Online Appendix. "Baseline" indicates the change predicted in the theoretical model from 1980 to 2015. Each of the remaining bars indicates 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 in 1980. "A" is total factor productivity. "LAT" is the set of occupation-specific efficiency indices.

Figure 3: Decomposition of the Change in Earnings Inequality from 1980 to 2015.

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 predicting a change in labor market inequality which is comparable in magnitude to what we observe in the data.

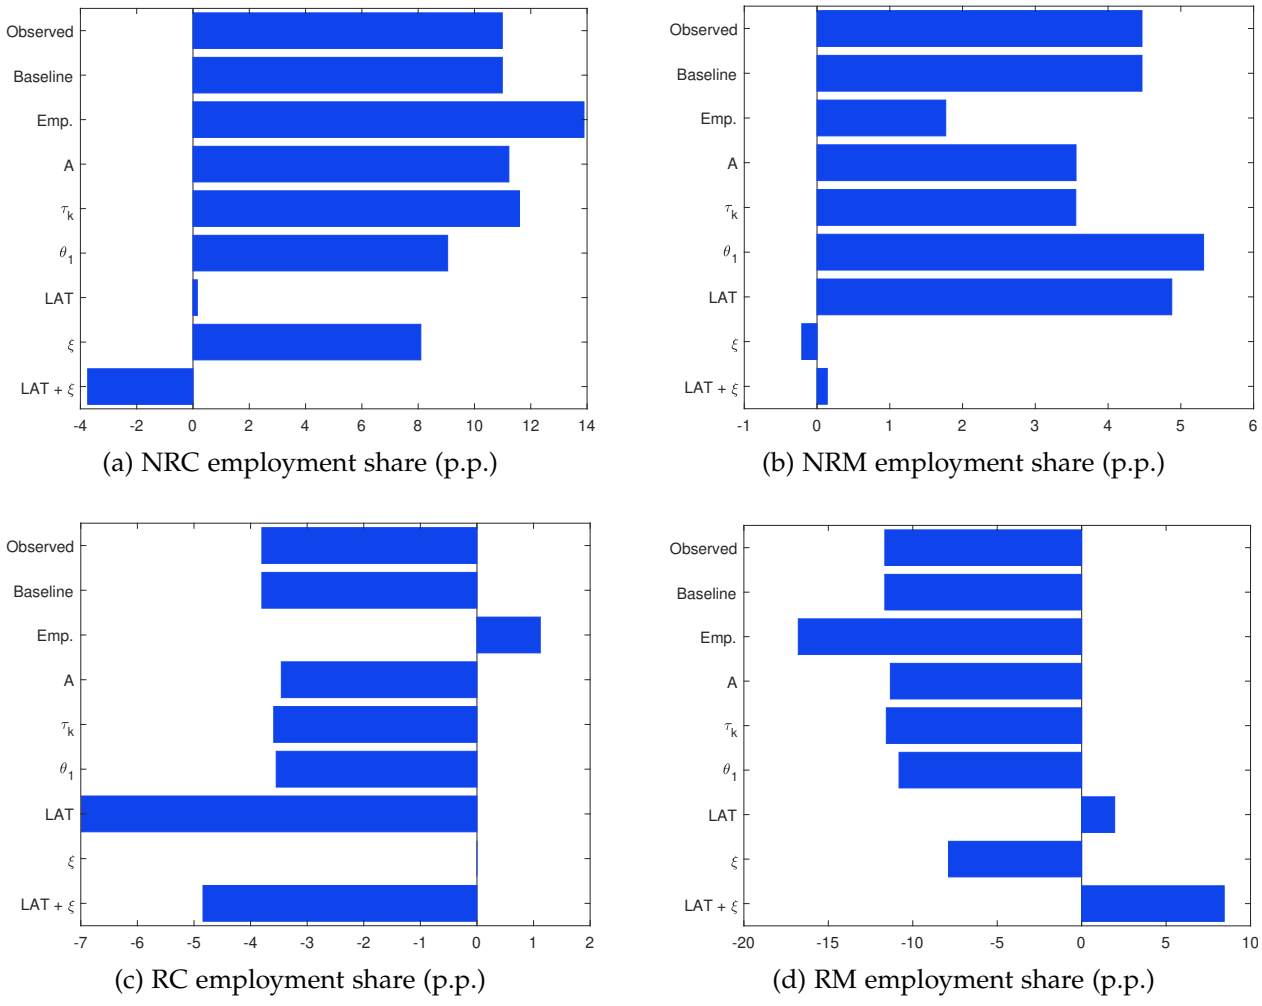
The drop in the relative price of an investment, ξ, is the single most important source of the increase in pre-tax earnings inequality. If this drop had not occurred, our framework predicts a rise in earnings dispersion of only 10%. Thus, investment-specific technological change alone accounts for two-thirds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

Latent occupation-biased technological change is next in terms of importance: 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 a lower increase in the NRC wage premium in particular. It only rises by 18 p.p.

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 are much less relevant. The reduction in tax progressivity had a positive but comparatively much smaller effect. Changes in the costs of acquiring the skills necessary to join each occupation while impacting the changes in employment shares (Figure 4) produced only a small increase in earnings inequality.



(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 3 for the description of the y-axis.

Figure 4: Decomposition of Variation in Employment Shares by Occupation Type from 1980 to 2015.

Figures 3b and 4 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. For 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 in negative territory in 1980 and increases by 6 p.p. (10 p.p. in the data). In other words, technological growth 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 by 50 p.p. and stands as one of the main sources of increased earnings dispersion in this model.

Concerning employment shares, the main takeaway is that technological change generates an increase 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 drop in 2015, and the share of NRM would remain unchanged.

In summary, technological change (and ISTC in particular) can 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 the next two subsections, we study optimal progressivity in 1980 and 2015 and discuss how technological change has affected it.

6.2 Optimal Tax Progressivity in 1980

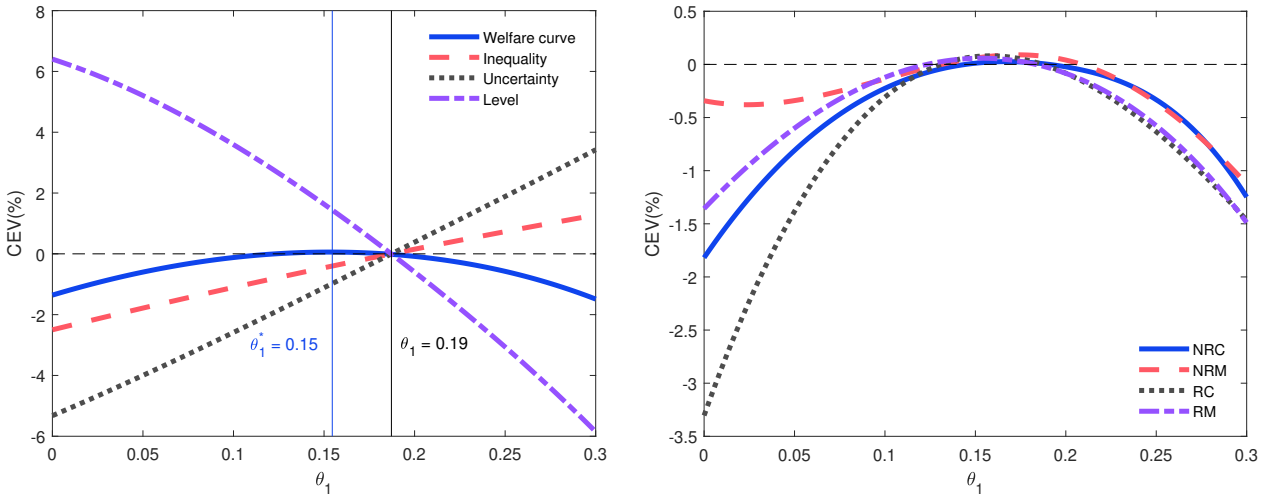
In this subsection, we examine the model-implied welfare function for the progressivity of the labor income tax system, θ_1. We decompose the welfare changes into the contributions from efficiency, redistribution and insurance (Flodén (2001)), and we study the welfare changes for the different occupations. Our main finding is that optimal progressivity is somewhat lower than the estimated progressivity of the U.S. labor income tax system in 1980 (0.15 vs. 0.19). Nonetheless, the model-implied gains from moving to the optimal level of θ_1 are very low at 0.06% in consumption equivalent variation.

We focus on a tax experiment where we keep the level of government spending exogenously fixed and focus on finding the most efficient way to cover the current spending level. This has the advantage that we do not have to make assumptions about the utility of government spending and has been a tradition in much of the literature on optimal taxation in Aiyagari-type OLG models[32]. For a given level of tax progressivity, θ_1, we adjust the tax level, θ_0, such that the government can raise enough tax revenue to cover the level of government expenditure, G, in the initial steady state. All other taxes are kept constant in our experiment.

To measure consumer welfare, we use the *utilitarian social welfare* criterion and maximize the expected lifetime utility of a household that is yet to enter the labor market

[32]See, e.g., Erosa and Gervais (2002), Conesa and Krueger (2006), Peterman (2016).

in a steady state.[33] 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 (insurance), (ii) the gain from reducing *inequality* in average lifetime marginal utilities of consumption and leisure (redistribution), and (iii) the impact that progressivity has on the overall *levels* of consumption and leisure via the incentive to work and invest (efficiency).



(a) Welfare Function

(b) Welfare Function by Occupation

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. 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 each year. In this case, welfare is the expected lifetime utility of agents once they have chosen occupation but before they know a_i (ability) and ϵ_{i1} (starting level of wage risk).

Figure 5: Social Welfare and Tax Progressivity in 1980

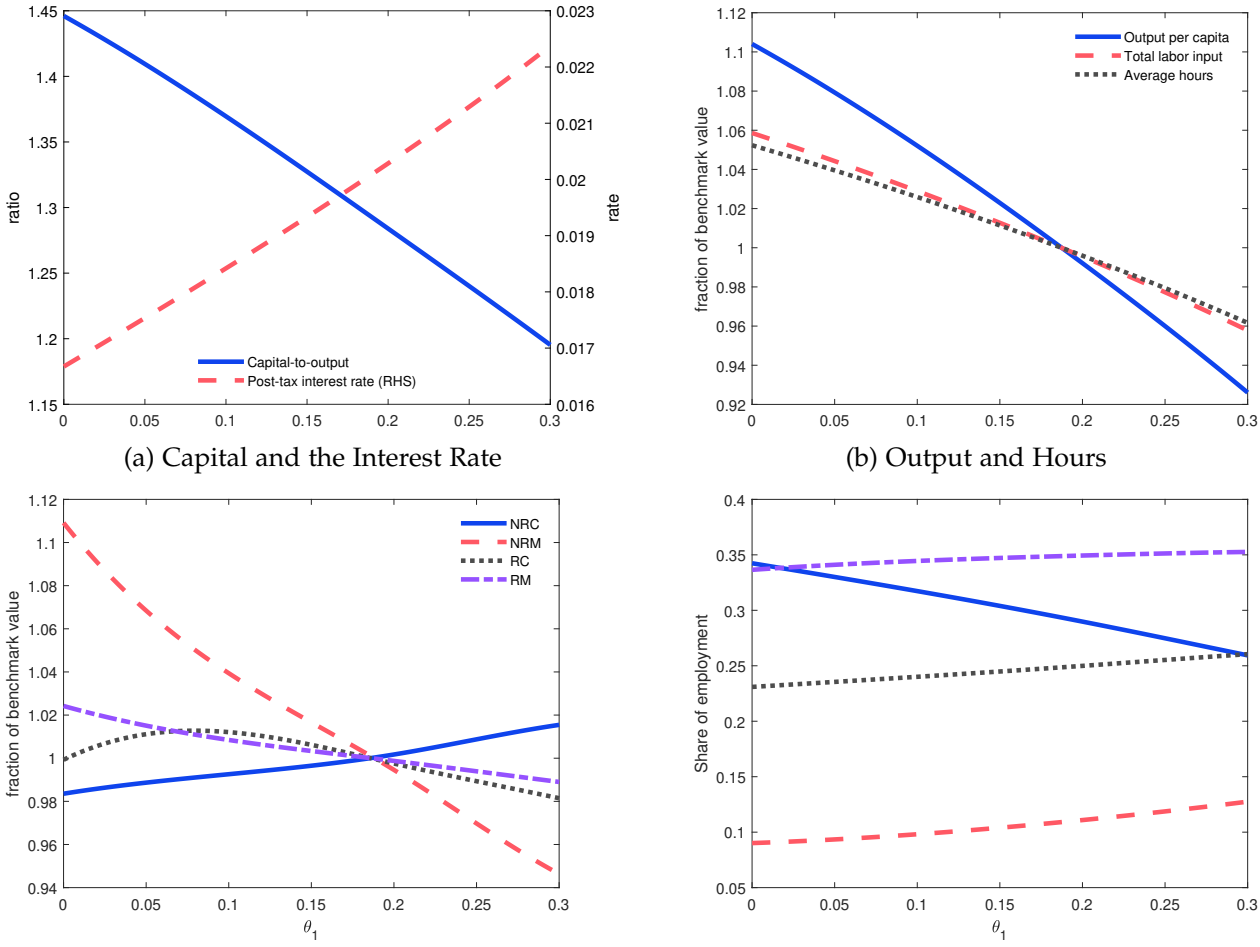
Figure 5 shows the results of our analysis. The left panel of Figure 5 plots the social welfare function for 1980 (the solid blue line). The model-implied optimal progressivity is 0.15, below the estimate of actual progressivity in 1980 of 0.19 by Wu (2021). Our estimate of optimal progressivity for 1980 is close to but below that of Heathcote et al. (2020), who put it at 0.18. Several differences between our models could lead to differences[34]. As we will see in Section 6.3 below, these differences become much more striking in 2015.

The low total welfare gain from moving to $\theta_1 = 0.15$ masks large countervailing

[33] See Online Appendix E for a definition of the welfare criterion and an exposition of how to decompose it into the contributions from different effects (efficiency, redistribution and insurance) in the case of our framework.

[34] We do for example have four occupations and capital-occupation complementarity. They have cross-sectional variation in the disutility of labor, which we do not.

movements in the contributions from different effects, which form the trade-offs faced by the policymaker. On the one hand, a reduction of θ_1 hampers the effectiveness of the labor income tax system as a mechanism to reduce both the uncertainty faced by the agents and the inequality between them. On the other hand, it increases the incentive to exert work effort[35], to save and to choose a higher-paying occupation (this also increases the marginal productivity in lower-paying professions). Figure 6 shows the comparative statics of output, wages and employment shares to progressivity in 1980. We look at the positive and negative effects of a reduction in progressivity in turn.



(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 6: Comparative Statics with Respect to Progressivity in 1980.

[35]For a given average tax rate, a larger θ_1 leads to a higher marginal tax rate and lower optimal choice of hours, see Holter et al. (2019)

On the plus side, reducing θ_1 to 0.15 increases aggregate saving and effort by reducing the marginal tax rate on high wage earners (in fact, most wage earners[36]). This raises the level of output per capita by 2% (Figure 6a). 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) how agents choose their occupation.

Figures 6c and 6d 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, the 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 productivity 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 the individuals in a given occupation do not necessarily coincide. This is most obvious for agents with a medium to the low cost of acquiring NRC training. These individuals benefit from an increase in progressivity, given that it discourages agents with a higher cost of NRC training from joining their occupation, reduces NRC labor input, and raises their marginal productivity and wages. This is the mechanism that explains why, in Figure 5b, agents who continue to choose NRC occupations when progressivity goes up prefer, on average, a high level of progressivity despite being at the top of the wage distribution.[37] Note that there is no selection effect in terms of the distribution of agents in the profession with respect to earnings potential when progressivity goes up. This is because the decision to join is taken only based on the training cost, and the occupation-specific ability has not yet been realized.

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

For middle-wage earners, in RC and RM occupations, wages increase by 0.6% and

[36] Holter et al. (2019) show that with this tax function, only very low earners will get an increase in their marginal tax rate when progressivity falls. The average tax rate will, however, increase for low-earners.

[37] This is the same mechanism that underlies the creation of professional guilds. One wishes to limit entry to the profession.

0.3%, and the employment shares drop by 0.3 p.p. and 0.4 p.p., respectively. In the end, the positive effects of increased output and wages for the lower-paid occupations dominate the negative effects. On the minus side, lowering progressivity reduces the insurance against idiosyncratic shocks by increasing the variance of after-tax labor income. It also reduces the ability of the tax system to reduce inequality between households.

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 barely. At the heart of the trade-off that determines optimal progressivity lies occupational choice. Whereas progressivity in most of the previous literature 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.

6.3 The Impact of Technological Change on Optimal Tax Progressivity

In this section, we answer the second main question in our paper: How did the technological transformation that took place between 1980 and 2015 affect optimal tax progressivity? To answer this, we evaluate optimal progressivity in 2015 and then investigate how each source of technological change, investment-specific technological change (ISTC), occupation-specific efficiency (LAT), and TFP affect our answer.

We find a significant drop in optimal tax progressivity, θ_1, from 0.15 to 0.05 between 1980 and 2015. The welfare gain from moving to the optimal policy in 2015 is 1.4% in consumption equivalent variation for unborn agents, whereas in 1980, it was only 0.06%. Technological change plays a major role in reducing optimal progressivity, and in particular, the single most important factor is ISTC. We begin by discussing the implications of changes in each source of technological change for aggregate variables. Then, we conduct a welfare analysis for 2015 and disentangle the drivers of the fall in optimal tax progressivity.

Table 5 displays the parameter changes that mimic the move to the 2015 calibration, and Table 7 displays how removing each source of technological change affects output per capita, capital, and prices. 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, by making an investment in equipment capital more attractive. This is reflected in a higher post-tax return rate when compared

Table 7: Impact of Technological Change on Quantities and Prices

	2015	No ISTC	1980 LAT	1980 TFP	No tech Δ
Output per 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: The table shows the equilibrium impact of technological change on quantities and prices in the model. "No ISTC', "1980 LAT', "1980 TFP' denote the values of aggregate variables when we, respectively, set investment prices at the 1980 value, the occupation-specific efficiency indices at 1980 values and TFP at its 1980 level, while keeping other parameters at their 2015 values. "No tech' shows the impact of simulataneously removing all sources of technological change. All variable (except the interest rate) are normalized to 1 in 1980.

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.

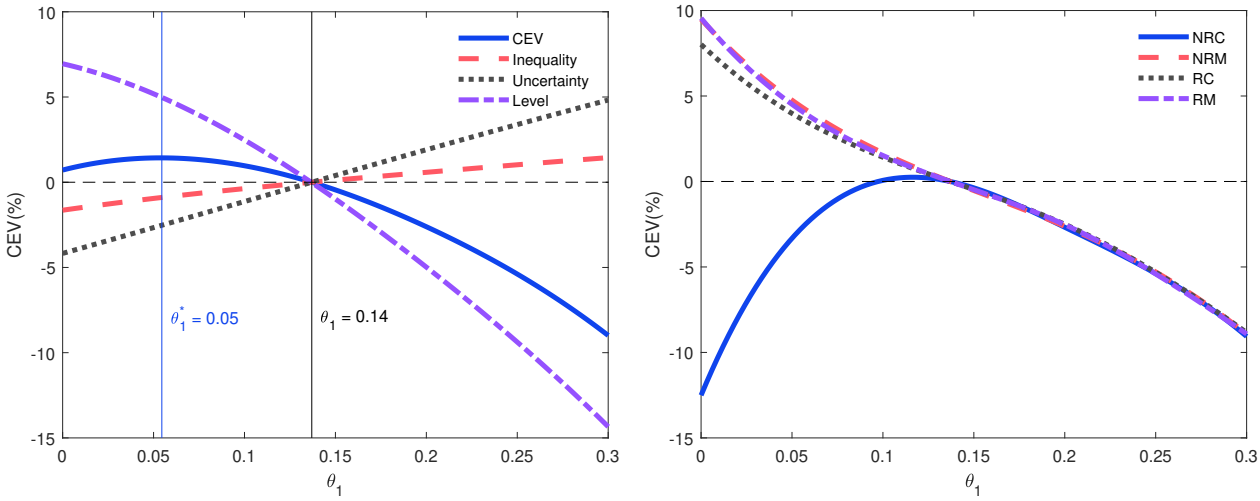
Latent occupation-biased technological change (LAT) has a positive but very small impact on GDP per capita and capital. However, it hurts the wage rate of all occupations except for NRC. This follows from the results of the production function estimation and the equilibrium occupation shares resulting from the changes in occupation-specific productivities. More workers end up choosing higher-paying occupations, thereby reducing their wage rates.

Finally, TFP has a large positive impact on GDP per capita, accounting for about 30% of its growth.[38] It is responsible for between 40 to 100% of the wage growth by occupation relative to the no tech growth scenario.

In summary, the technological change produced a significant increase in the level of output per capita, the returns to investment (interest rates), and wages. However, as we argued in section 6.1 it also generated an increase in labor market inequality, with the NRC wages growing more than those of other occupations. This affects the trade-offs between efficiency, inequality and insurance that underlie the determination of optimal tax progressivity. Figure 7 displays the welfare analysis.

Wu (2021) estimates that the progressivity in the U.S. tax system declined from 0.19 to 0.14 between 1980 and 2015. He finds that changes in economic conditions can explain

[38]The decompositions need not add to 100% as they interact.



(a) Welfare Function

(b) Welfare Function by Occupation

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. 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 each year. In this case, welfare is the expected lifetime utility of agents once they have chosen occupation but before they know a_i (ability) and ϵ_{i1} (starting level of wage risk).

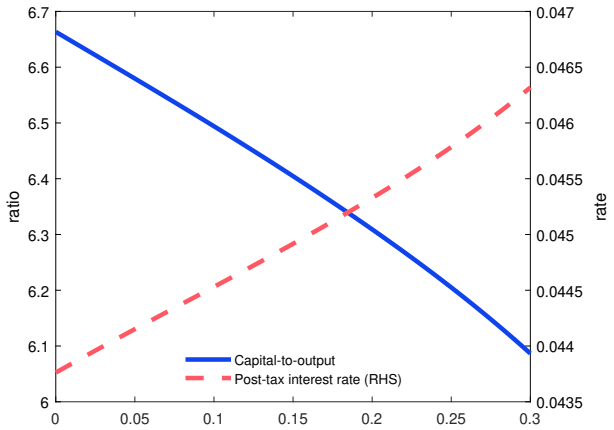
Figure 7: Optimal Progressivity in 2015, for all and Across Occupations

about 62% of this change[39], and argues that the rest could be due to the government shifting its welfare weights towards high-ability households. In our framework, we do, however, predict that optimal progressivity dropped from 0.15 in 1980 to 0.05 in 2015 without changing welfare weights.

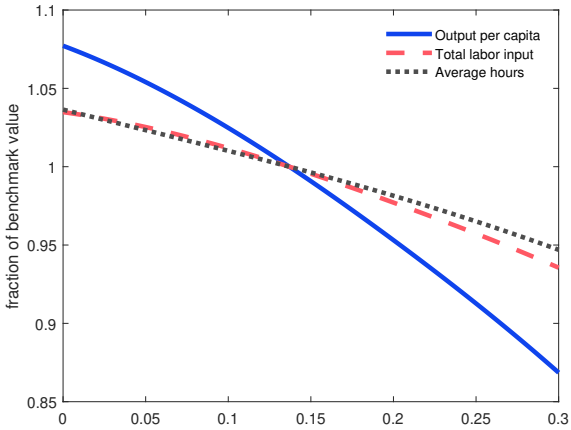
What leads to this remarkable two-thirds drop in optimal progressivity? How do the different sources of technological transformation contribute to this shift in optimal progressivity, and how do they affect the trade-offs faced by the policymaker? To answer these questions, we use our decomposition of the welfare function into the level, inequality and insurance effects and study the impact of each source of technological change (Figures 9 and 10). In a nutshell, we find that ISTC is the main driver of the results.

The trade-off between efficiency and redistribution/insurance is now more significantly dominated by the efficiency side (Figures 7a and 10). This is ultimately a quantitative outcome due to the estimated production function and changes in technology

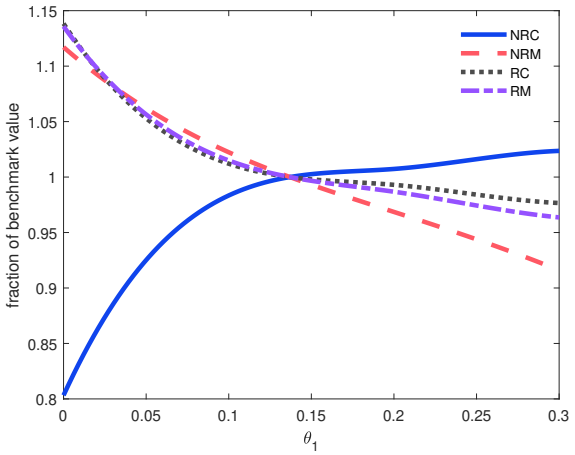
[39] An ageing population and shrinking gender wage gap calls for less progressive taxes, increased idiosyncratic risk calls for more progressive taxes, and an increase in the skill premium (modeled with a parameter governing the returns to human capital investment) is about neutral with respect to optimal tax progressivity.



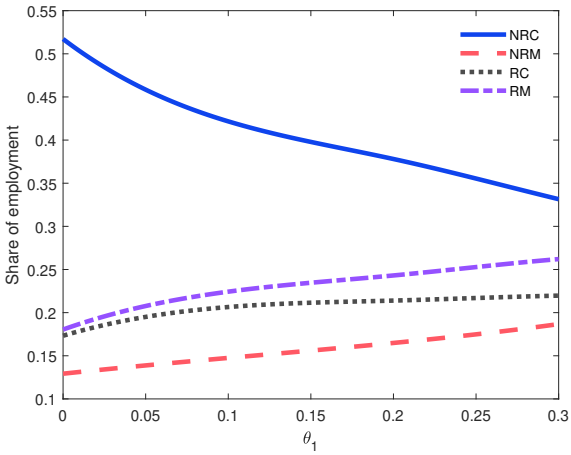
(a) Capital and the Return Rate



(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. Panels show how prices and quantities in the economy change concerning progressivity.

Figure 8: Comparative Statics of Progressivity in 2015.

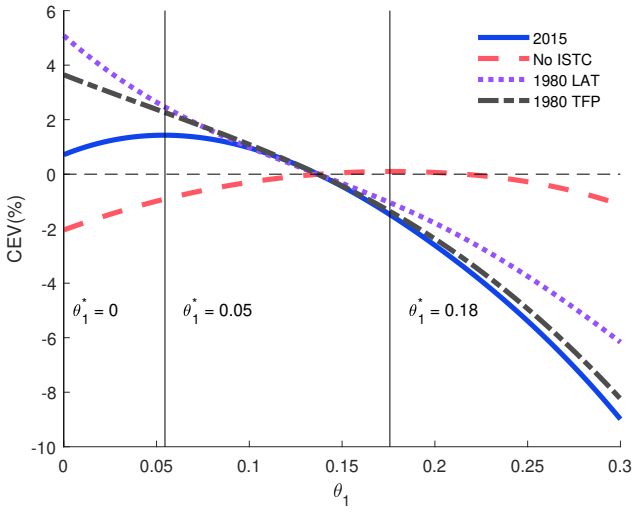
over time. However, it is due to three main reasons: (i) There is now much more capital, making NRC workers very productive. Getting more workers to choose NRC has a stronger positive impact on output, the *level* channel. (ii) The higher equilibrium return rate on capital given its high productivity. This implies that self-insurance is easier and the welfare gains from improving risk-sharing through more progressive taxation are less significant, which reduces the importance of the *uncertainty* channel relative to 1980. (iii) The effect of capital accumulation and inflow of workers to NRC professions during the 1980-2015 period on the marginal productivities and employment shares in the lower paid professions is such that their wage rates increase substantially. This dampens the *inequality* channel.

There is, on the one hand, still a positive welfare effect from reducing inequality through progressivity. 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, this channel is flattened because employment shares change in equilibrium as agents select different occupations in response to variation in relative wages and post-tax earnings (reason iii). 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 8 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 positive impact of a reduction in progressivity on the after-tax earnings of NRC occupations, which are the highest-paying, makes them more attractive. This 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 displayed in Figure 7b. Are the agents in NRC occupations partial to a reduction in progressivity (conditional on still choosing NRC)? 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 tax rate on high wage earners increases their after-tax earnings but leads to an inflow of workers and actually a reduction in their hourly post-tax wages at some point[40]. As progressivity drops, 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 insurance against uncertainty. However, as previously discussed, the higher return rate on savings relative to the 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

[40]That 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.

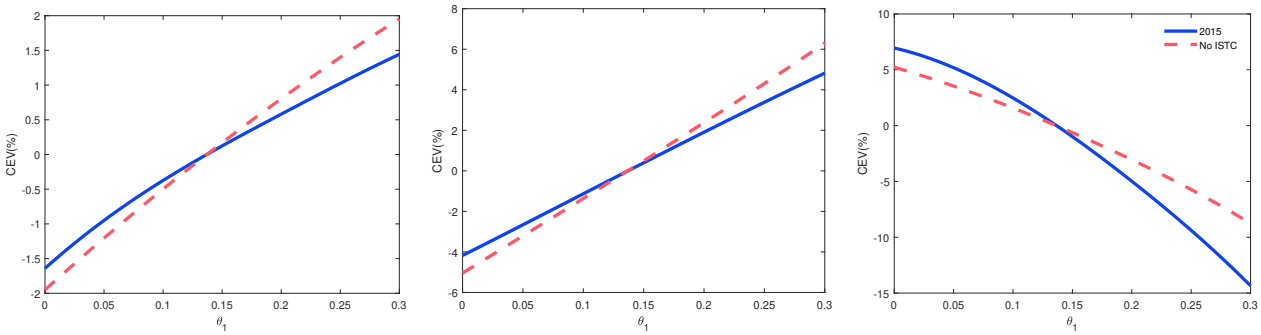
Figure 9: Optimal Progressivity in 2015 and the Role of Technological Change.

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.

Quantitatively, these results stand in stark contrast to those obtained by Heathcote et al. (2020), who find an optimal level of $\theta_1 = 0.16$ in 2016 (a slight drop from 1980)[41]. Part of the reason for this is the heterogeneous impact of savings and technological change on the marginal productivities of each occupation, which is absent from their model. Rather, in their framework, agents choose their level of skills, which are imperfect substitutes in production but have a single constant elasticity of substitution, and there is no role for capital in production.

Figure 9 confirms our interpretation of ISTC as the main driver of these changes. The solid blue line is the welfare function in 2015, where the maximum is 0.05. The dashed red line indicates the calibration for 2015 with equipment investment prices at their 1980 level. Here, the maximum is 0.18, close to the value estimated for 1980 and the optimal θ_1 found by Heathcote et al. (2020) and Wu (2021). As discussed earlier, this equilibrium

[41] Wu (2021) also finds an optimal value of $\theta_1 = 0.16$ in the 2010s.



(a) Inequality Channel (b) Uncertainty Channel (c) Level Channel

Note: Each graph shows the impact of changing progressivity on each element of the welfare function. "2015" indicates the breakdown for the 2015 calibration. "No ISTC" shows the breakdown for the 2015 calibration without investment-specific technological change.

Figure 10: Welfare Trade-offs With and Without ISTC.

displays lower return rates and 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, earnings 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 growth raises wages and output. It does, however, have little effect on relative wages, see Table 7. When everyone becomes poorer, this, however, leads to fewer savings and capital, which moves optimal progressivity to the left.

Figure 10 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 same breakdown 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 savings by households as a result of lower progressivity.

7 Conclusion

We have developed a life-cycle, overlapping generations model with uninsurable idiosyncratic earnings risk, three sources of technological change, a detailed tax system, and occupational choice. Furthermore we have estimated an aggregate production function with capital-occupation complementarity and four types of labor inputs that differ with respect to cognitive complexity and routine task intensity. Inserting the estimated production technology in our model, calibrating it to resemble the U.S. in 1980, and then treating it with the changes in the price of equipment goods (ISTC), latent occupation-biased technological change (LAT) and TFP, we have shown that technological transformation can fully account for the change in wage premia as well as the increase in earnings inequality between 1980 and 2015. The main driver is Investment-Specific Technological Change which leads to more capital accumulation, increasing the 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 a significant drop in optimal tax progressivity between 1980 and 2015. This fall can be solely attributed to ISTC. In our model, in addition to the traditional effects of increasing work hours and savings, lower progressivity leads to an inflow of workers into higher-paid occupations, which are more productive with higher ISTC. This raises output and also the wages of those remaining in the occupations at the bottom of the wage distribution, dampening the redistributive benefits of progressive taxation. Finally, ISTC raises the real return rates on saving, making self-insurance easier and thus weakening the insurance role of progressive taxation.

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 or training to enter better-paid occupations, see e.g. Krueger and Ludwig (2016), Stantcheva (2018). There will be interesting interactions between these policies, the tax system, occupational choice and wages. Second, we did not study capital or wealth taxation in this paper. However, the importance of capital-occupation complementarities demonstrated in this paper could likely alter conclusions on optimal capital taxation. Finally, we do not consider 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?

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. (2018). Artificial Intelligence, Automation and Work. Boston University - Department of Economics - Working Papers Series dp-298, Boston University - Department of Economics.

Acikgoz, O., Hagedorn, M., Holter, H., and Wang, Y. (2022). The Optimum Quantity of Capital and Debt. Working paper.

Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, 109(3):659–684.

Ales, L., Kurnaz, M., and Sleet, C. (2015). Technical change, wage inequality, and taxes. *American Economic Review*, 105(10):3061–3101.

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.

Badel, A. and Huggett, M. (2017). The sufficient statistic approach: Predicting the top of the laffer curve. *Journal of Monetary Economics*, 87:1–12.

Bakis, O., Kaymak, B., and Poschke, M. (2015). Transitional dynamics and the optimal progressivity of income redistribution. *Review of Economic Dynamics*, 18(3):679–693.

Barro, R. (2000). Inequality and growth in a panel of countries. *Journal of Economic Growth*, 5(1):5–32.

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.

Boar, C. and Midrigan, V. (2022). Efficient Redistribution. 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.

Brinca, P., Homem Ferreira, M., Franco, F. A., Holter, H. A., and Malafry, L. (2021). Fiscal consolidation programs and income inequality. *International Economic Review*, 62(1):405–460.

Cameron, A. and Trivedi, P. (2005). *Microeconometrics*. Cambridge University Press.

Chakraborty, I., Holter, H. A., Hai, R., and Stepanchuk, S. (2017). The real effects of financial (dis)integration: A multi-country equilibrium analysis of europe. *Journal of Monetary Economics*, (85):28–45.

Conesa, J. C., Kitao, S., and Krueger, D. (2009). Taxing capital? not a bad idea after all! *American Economic Review*, 99(1):25–48.

Conesa, J. C. and Krueger, D. (2006). On the optimal progressivity of the income tax code. *Journal of Monetary Economics*, 53:1425–1450.

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*.

Dyrda, S. and Pedroni, M. (2021). Optimal Fiscal Policy in a Model with Uninsurable Idiosyncratic Shocks. *Working Paper*.

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

Erosa, A. and Gervais, M. (2002). Optimal taxation in life-cycle economies. *Journal of Economic Theory*, 105(2):338–369.

Ferriere, A., Navarro, G., Board, F. R., and Vardishvili, O. (2022). On the Optimal Design of Transfers and Income-tax Progressivity.

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. (2022). Should Robots Be Taxed? *The Review of Economic Studies, forthcoming*.

Guvenen, F., Kuruscu, B., and Ozkan, S. (2014). Taxation of human capital and wage inequality: A cross-country analysis. *Review of Economic Studies*, 81:818–850.

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. (2020). Presidential address 2019: How should tax progressivity respond to rising income inequality? *Journal of the European Economic Association*, 18(6):2715–2754.

Heathcote, J., Storesletten, S., and Violante, G. (2017). Optimal tax progressivity: An analytical framework. *Quarterly Journal of Economics*, 134:1693–1754.

Herrington, C. (2015). "public education financing, earnings inequality, and intergenerational mobility. *Review of Economic Dynamics*, 18:822–842.

Holter, H. (2015). Accounting for cross-country differences in intergenerational earnings persistence: The impact of taxation and public education expenditure. *Quantitative Economics*, 6(2):385–428.

Holter, H. A., Krueger, D., and Stepanchuk, S. (2019). How do tax progressivity and household heterogeneity affect laffer curves? *Quantitative Economics*, 10(4):1317–1356.

Hugget, M. (1993). The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies. *Journal of Economic Dynamics and Control*, 17:953–969.

Huggett, M., Ventura, G., and Yaron, A. (2011). Sources of lifetime inequality. *American Economic Review*, 101(7):2923–2954.

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.

Kaplan, G. and Zoch, P. ((2020)). Markups, Labor Market Inequality and the Nature of Work. Working paper, University of Chicago.

Karabarbounis, L. and Neiman, B. (2014). The global decline of the labor share. *The Quarterly Journal of Economics*, 129(1):61–103.

Kindermann, F. and Krueger, D. (2022). High marginal tax rates on the top 1%? lessons from a life cycle model with idiosyncratic income risk. *American Economic Journal: Macroeconomics, forthcoming*, 14(2):319–366.

Krueger, D. and Ludwig, A. (2016). On the optimal provision of social insurance: Progressive taxation versus education subsidies in general equilibrium. *Journal of Monetary Economics*, 77:72–98.

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.

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.

Peterman, W. (2016). The effect of endogenous human capital accumulation on optimal taxation. *Review of Economic Dynamics*, 21:46–71.

Scheuer, F. and Werning, I. (2015). The taxation of superstars. *Quarterly Journal of Economics*, 132(1):211–270.

Stantcheva, S. (2018). Optimal taxation and human capital policies over the life cycle. *Journal of Political Economy*, 125(6):1931–1990.

Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economic Letters*, 20:177–181.

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

Wu, C. (2021). More unequal income but less progressive taxation. *Journal of Monetary Economics*, 117(C):949–968.

ONLINE APPENDIX

The Online 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 our welfare measures and how to decompose welfare effect of policy changes into the effect on redistribution, insurance and efficiency. Section F derives the first order conditions of the firms in our model. In Section G, we study optimal tax policy in 2015 under the alternative assumption of fixed employment shares (i.e. agents have no choice of occupation).

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 *INCWAGE*, *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. The population of interest comprises non-military and non-institutionalized individuals aged 16 to 70, excluding the self-employed and farm sector workers. 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,

Table 8: CPS Sample Selection (survey years 1968-2017)

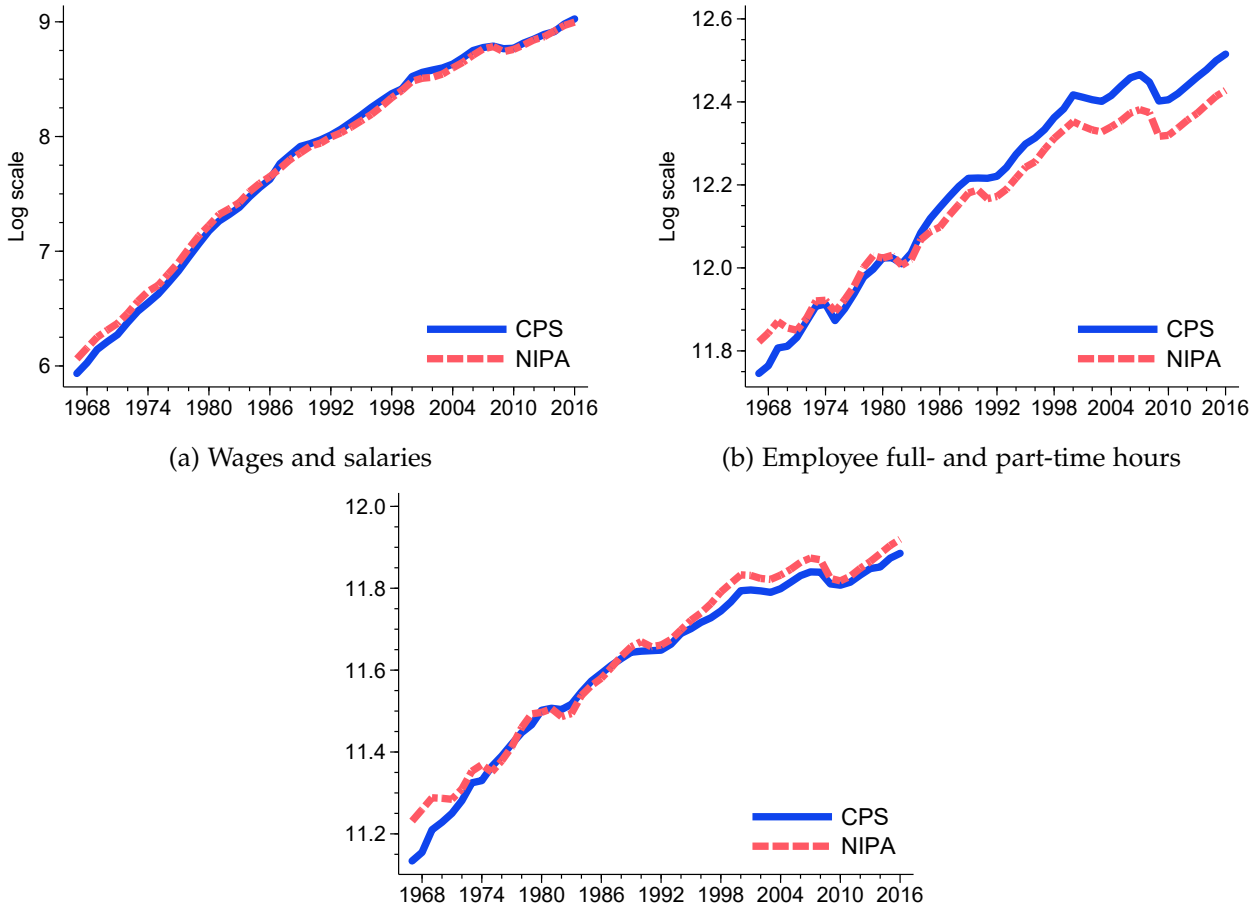
	Dropped	Remaining
Initial sample		**4,089,617**
Wage > 0.5 × min. wage	116,608	3,973,009
Sample A		**3,973,009**
Age 25-64	861,598	3,111,411
Hours worked per week last year > 6	19,308	3,092,103
Sample B		**3,092,103**

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 validate the data, we compare a set 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) 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.

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: n over-sample of roughly 2000 poor families selected from the Survey of Economic Opportunities (SEO), 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 × 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 A, 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 $l_{it} = h_{it} w k_{it}$,

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 $(w_{\text{NRC},t}, w_{\text{NRM},t}, w_{\text{RC},t}, w_{\text{RM},t})$.

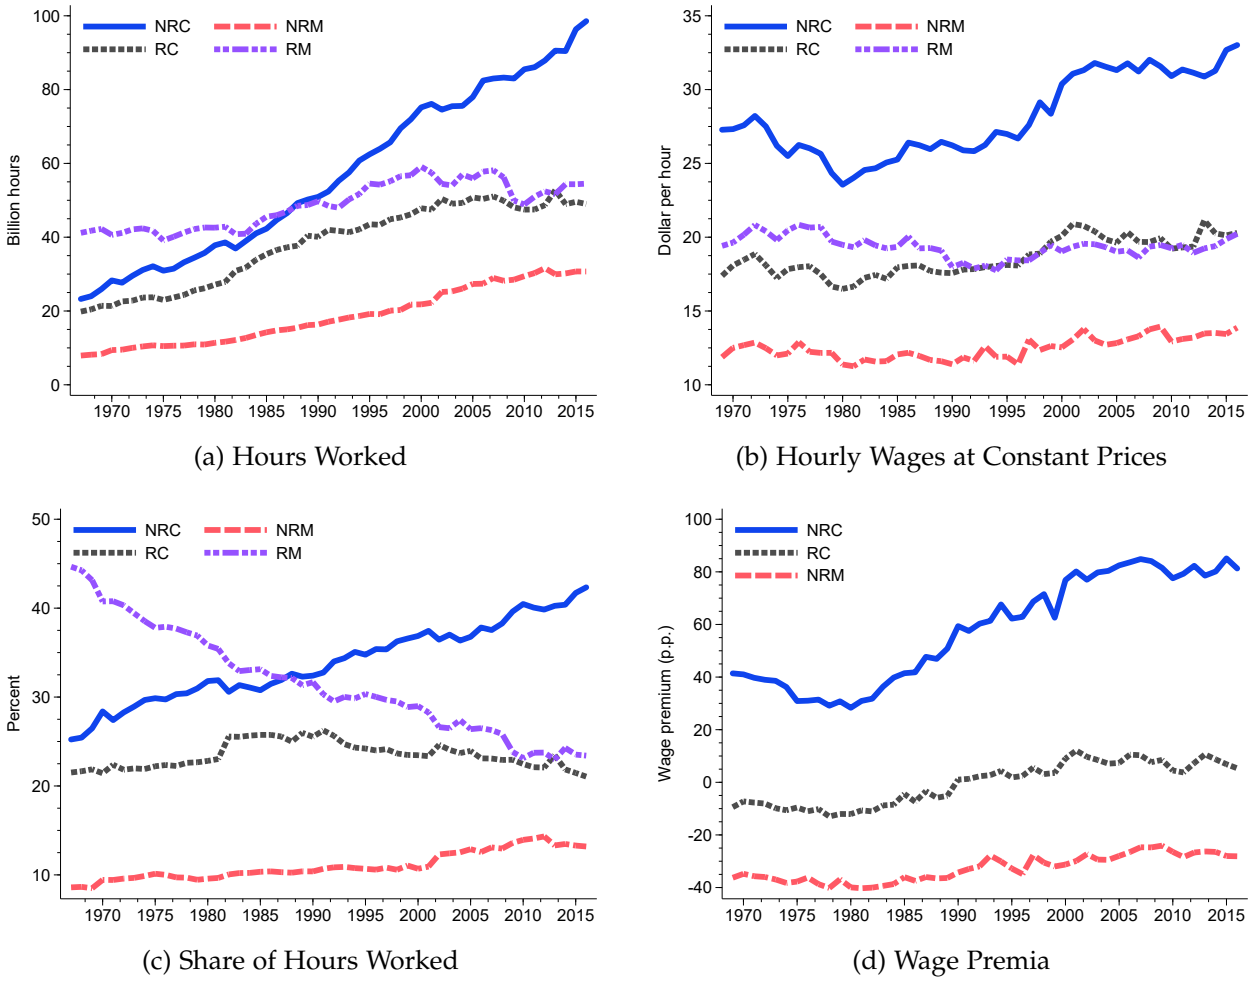
B.2 Capital, Prices and Output

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



(a) Hours Worked

(b) Hourly Wages at Constant Prices

(c) Share of Hours Worked

(d) Wage Premia

Note: Wage premia are obtained as the log difference between the constant composition average wage of each occupation category. Groups for wages are constructed by using a constant composition of individual observable characteristics. The data source is the CPS Annual Social and Economic Supplement. See sections A and B of the Online Appendix for details.

Figure B.1: Employment and Wages by Occupation Category.

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, $P_{it}\text{NetStock}_{it}$, and depreciation at current cost, $P_{it}\text{Dep}_{it}$, to compute depreciation rates, which are given by the following

formula:

$$\delta_{it} = \frac{P_{it}\text{Dep}_{it}}{P_{it}\text{NetStock}_{it} + P_{it}\text{Dep}_{it}}.$$

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:

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

$f(.)$ 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 structure 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:

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

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 21 using a bisection method. Compute marginal productivities A-24-A-28 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 structure 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 structure 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.[42]

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.

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

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:

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

Utilitarian social welfare is defined as:

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

where $\Phi(s)$ is the distribution over the idiosyncratic occupation costs after occupation decisions are made. The problem solved by our social planner is:

$$\max_{\{\theta_1, \theta_0\}} W, \quad s.t. \quad G^* = \int \tau_k r(b + \Gamma) + \tau_c c + n\tau_l \left[\frac{nw(j, o, a, u)}{1 + \tilde{\tau}_{ss}} \right] d\Phi \tag{A-3}$$

That is, the social planner takes government spending as given and find the socially optimal progressivity and level of the tax system to raise the required revenue. 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:

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

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

$$v(s; g(s)) = \kappa(s) + \sum_{j=1}^{J} \beta^{j-1} S_j \log(1 + g(s)) + \mathbb{E}_0 \left[\sum_{j=1}^{J} \beta^{j-1} \left[S_j u(c_j, n_j) + (1 - S_j) D(b_{j+1}) \right] \right]. \tag{A-5}$$

Integrating over the idiosyncratic state s:

$$W(g) = \log(1 + g) \sum_{j=1}^{J} \beta^{j-1} S_j + W. \tag{A-6}$$

Let $v(s, B)$ denote the expected lifetime utility of a household with starting state s under an alternative government 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:

$$v(s; g_U(s)) = v(s, B), \tag{A-7}$$

$$W(g_U) = W(B). \tag{A-8}$$

Given these definitions we can define the consumption equivalent variation for an unborn individual by substituting A-6 into A-8:

$$\log(1 + g_U) \sum_{j=1}^{J} \beta^{j-1} S_j + W = W(B). \tag{A-9}$$

Solving for g_U:

$$g_U = \exp\left(\frac{W(B) - W}{\sum_{j=1}^{J} \beta^{j-1} S_j}\right) - 1. \tag{A-10}$$

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:

$$\sum_{j=1}^{J} u(\bar{c}, \bar{n}) = v(s), \tag{A-11}$$

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-11 for $\bar{c}$:

$$\bar{c} = \exp\left(\frac{v(s)}{\sum_{j=1}^{J} \beta^{j-1}} + \chi \frac{\bar{n}^{1+\eta}}{1+\eta}\right). \tag{A-12}$$

We can now define the *cost of inequality*:

$$\sum_{j=1}^{J} \beta^{j-1} u((1-\rho_{\text{ine}})\bar{C}, \bar{N}) = \int \sum_{j=1}^{J} u(\bar{c}, \bar{n})\, d\Phi(s) = W, \tag{A-13}$$

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}:

$$\rho_{\text{ine}} = 1 - \exp\left(\frac{W}{\sum_{j=1}^{J} \beta^{j-1}} - u(\bar{C}, \bar{N})\right). \tag{A-14}$$

The *cost of uncertainty* is defined as:

$$\sum_{j=1}^{J} \beta^{j-1} u((1-\rho_{\text{unc}})C, N) = \sum_{j=1}^{J} \beta^{j-1} u(\bar{C}, \bar{N}), \tag{A-15}$$

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

$$\rho_{\text{unc}} = 1 - \exp\left(u(\bar{C}, \bar{N}) - u(C, N)\right). \tag{A-16}$$

Finally, we are ready to asses 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 $\tilde{C}^B$:

$$\sum_{j=1}^{J} \beta^{j-1} u(\tilde{C}^B, N^A) = \sum_{j=1}^{J} \beta^{j-1} u(C^B, N^B). \tag{A-17}$$

Solving for $\tilde{C}^B$:

$$\tilde{C}^B = \exp\left(\log C^B - \chi \frac{N^{B^{1+\eta}} - N^{A^{1+\eta}}}{1+\eta}\right). \tag{A-18}$$

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:

$$g_{\text{lev}} = \frac{\tilde{C}^B}{C^A} - 1. \tag{A-19}$$

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

$$g_{\text{ine}} = \frac{1-\rho_{\text{ine}}^{B}}{1-\rho_{\text{ine}}^{A}} - 1. \tag{A-20}$$

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

$$g_{\text{unc}} = \frac{1-\rho_{\text{unc}}^{B}}{1-\rho_{\text{unc}}^{A}} - 1. \tag{A-21}$$

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

$$g_{U} = (1+g_{\text{lev}})(1+g_{\text{ine}})(1+g_{\text{unc}}) - 1. \tag{A-22}$$

F First-Order Conditions of the Firms

The first-order conditions of the firms are the following:[43]

$$w_{\text{NRC},t} = \Xi_t \varphi_1 \left[\phi_1 \left(\frac{K_{e,t}}{N_{\text{NRC},t}} \right)^{\frac{\rho_1 - 1}{\rho_1}} + (1-\phi_1) \right]^{\frac{\sigma-\rho_1}{(\rho_1-1)\sigma}} [1-\phi_1] \varrho_{\text{NRC},t}, \tag{A-23}$$

$$w_{\text{NRM},t} = \Xi_t \varphi_2 \left[\phi_2 \left(\frac{K_{e,t}}{N_{\text{NRC},t}} \right)^{\frac{\rho_2 - 1}{\rho_2}} + (1-\phi_2) \left(\frac{N_{\text{NRM},t}}{N_{\text{NRC},t}} \right)^{\frac{\rho_2 - 1}{\rho_2}} \right]^{\frac{\sigma-\rho_2}{(\rho_2-1)\sigma}} [1-\phi_2] \left(\frac{N_{\text{NRM},t}}{N_{\text{NRC},t}} \right)^{-\frac{1}{\rho_2}} \varrho_{\text{NRM},t}, \tag{A-24}$$

$$w_{\text{RC},t} = \Xi_t \varphi_3 \left[\phi_3 \left(\frac{K_{s,t}}{N_{\text{NRC},t}} \right)^{\frac{\rho_3 - 1}{\rho_3}} + (1-\phi_3) \left(\frac{N_{\text{RC},t}}{N_{\text{NRC},t}} \right)^{\frac{\rho_3 - 1}{\rho_3}} \right]^{\frac{\sigma-\rho_3}{(\rho_3-1)\sigma}} [1-\phi_3] \left(\frac{N_{\text{RC},t}}{N_{\text{NRC},t}} \right)^{-\frac{1}{\rho_3}} \varrho_{\text{RC},t}, \tag{A-25}$$

$$w_{\text{RM},t} = \Xi_t (1-\varphi_1-\varphi_2-\varphi_3) \left(\frac{N_{\text{RM},t}}{N_{\text{NRC},t}} \right)^{-\frac{1}{\sigma}} \varrho_{\text{RM},t}, \tag{A-26}$$

[43] Marginal 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.

$$r_{s,t} = A_t \alpha \left[\frac{K_{e,t}}{N_{\text{NRC},t}}\right]^{\alpha-1} \Lambda_t^{\frac{\sigma(1-\alpha)}{\sigma-1}}, \tag{A-27}$$

$$\begin{aligned} r_{e,t} = \Xi_t \Bigg[& \varphi_1 \left(\phi_1 \left[\frac{K_{e,t}}{N_{\text{NRC},t}}\right]^{\frac{\rho_1-1}{\rho_1}} + [1-\phi_1] \right)^{\frac{\sigma-\rho_1}{(\rho_1-1)\sigma}} \phi_1 \left(\frac{K_{e,t}}{N_{\text{NRC},t}}\right)^{-\frac{1}{\rho_1}} + \\ & \varphi_2 \left(\phi_2 \left[\frac{K_{e,t}}{N_{\text{NRC},t}}\right]^{\frac{\rho_2-1}{\rho_2}} + [1-\phi_2] \left[\frac{N_{\text{NRM},t}}{N_{\text{NRC},t}}\right]^{\frac{\rho_2-1}{\rho_2}} \right)^{\frac{\sigma-\rho_2}{(\rho_2-1)\sigma}} \phi_2 \left(\frac{K_{e,t}}{N_{\text{NRC},t}}\right)^{-\frac{1}{\rho_2}} + \\ & \varphi_3 \left(\phi_3 \left[\frac{K_{e,t}}{N_{\text{NRC},t}}\right]^{\frac{\rho_3-1}{\rho_3}} + [1-\phi_3] \left[\frac{N_{\text{RC},t}}{N_{\text{NRC},t}}\right]^{\frac{\rho_3-1}{\rho_3}} \right)^{\frac{\sigma-\rho_3}{(\rho_3-1)\sigma}} \phi_3 \left(\frac{K_{e,t}}{N_{\text{NRC},t}}\right)^{-\frac{1}{\rho_3}} \Bigg], \end{aligned} \tag{A-28}$$

where[44]

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

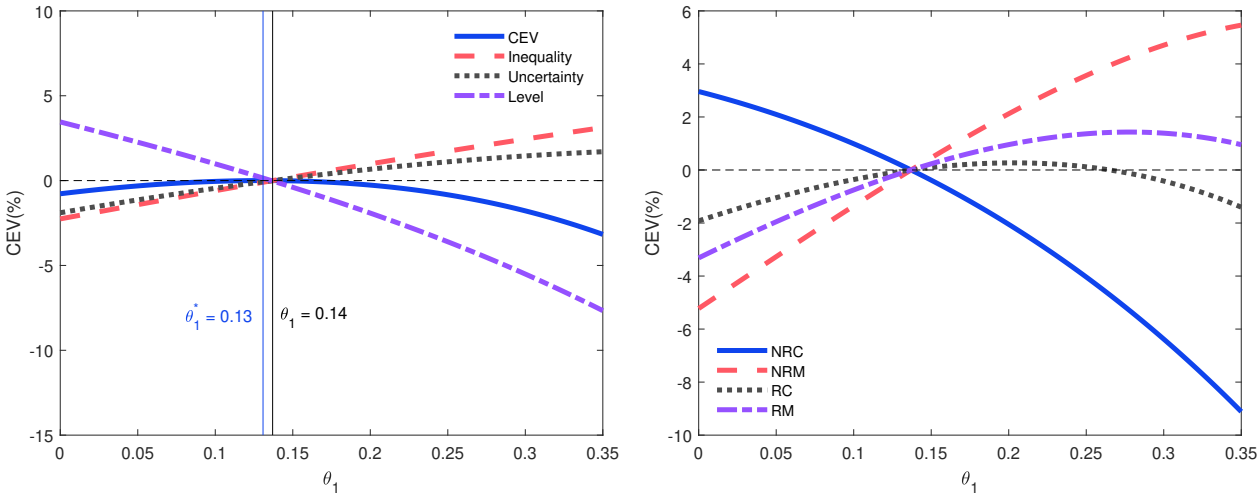
G The Impact of Occupational Choice

How would our results in Section 6.3 change if we ignored occupational choice and the equilibrium effects it has on wages and other variables? Figure G.1 shows how abstracting from occupational choice, and assuming fixed 2015 employment shares, 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 steepness of the level channel remains practically unchanged compared to the economy with occupational choice.

[44] The variable Λ_t is defined as:

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



(a) Welfare Function

(b) Welfare Function by Occupation

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 been assigned 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.

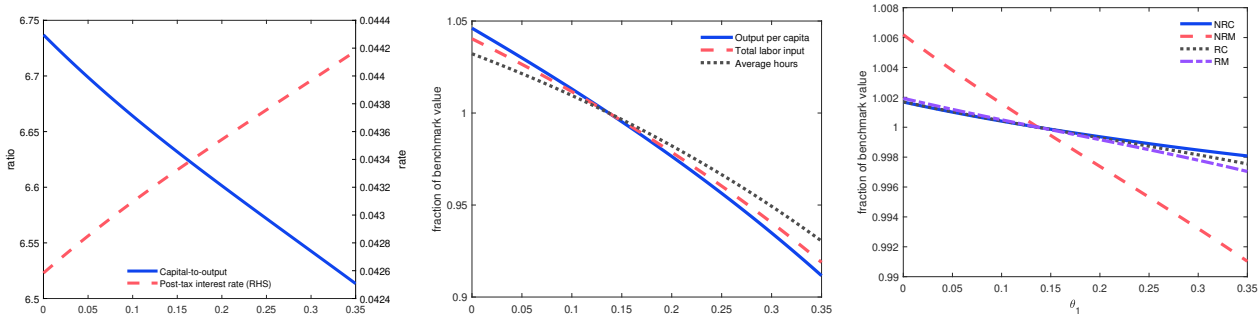
Figure G.1: Optimal Progressivity in 2015 and Across Occupations, Without Occupational Choice.

The key difference in this experiment is the effects of progressivity on occupational choice and wages. 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 G.2).

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-tax earnings, which point to a greater strength of the inequality channel. This analysis is illustrated by the breakdown of the welfare gain by occupation displayed in Figure G.1b 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.



(a) Capital and the Interest Rate (b) Output and Hours (c) Wages 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. Panels show how prices and quantities in the economy change with respect to progressivity.

Figure G.2: Comparative Statics With Respect to Tax Progressivity in 2015 in the Absence of Occupational Choice.