INTERNATIONAL MONETARY FUND

AI Adoption and Inequality

Emma Rockall, Marina M. Tavares, and Carlo Pizzinelli

WP/25/68

2025
APR

WORKING PAPER

 WP/25/68

IMF Working Paper
Research Department

AI Adoption and Inequality
Prepared by Emma Rockall, Marina M. Tavares, and Carlo Pizzinelli

Authorized for distribution by Florence Jaumotte
April 2025

ABSTRACT: There are competing narratives about artificial intelligence's impact on inequality. Some argue AI will exacerbate economic disparities, while others suggest it could reduce inequality by primarily disrupting high-income jobs. Using household microdata and a calibrated task-based model, we show these narratives reflect different channels through which AI affects the economy. Unlike previous waves of automation that increased both wage and wealth inequality, AI could reduce wage inequality through the displacement of high-income workers. However, two factors may counter this effect: these workers' tasks appear highly complementary with AI, potentially increasing their productivity, and they are better positioned to benefit from higher capital returns. When firms can choose how much AI to adopt, the wealth inequality effect is particularly pronounced, as the potential cost savings from automating high-wage tasks drive significantly higher adoption rates. Models that ignore this adoption decision risk understating the trade-off policymakers face between inequality and efficiency.

JEL Classification Numbers:	J21, O33, D31
Keywords:	Artificial intelligence, Employment, Inequality
Author's E-Mail Address:	erockall@stanford.edu, mmendestavares@imf.org, cpizzinelli@imf.org

AI Adoption and Inequality*

Emma Rockall[†], Marina M. Tavares[‡] and Carlo Pizzinelli[‡]

February 2025

Abstract

There are competing narratives about artificial intelligence's impact on inequality. Some argue AI will exacerbate economic disparities, while others suggest it could reduce inequality by primarily disrupting high-income jobs. Using household microdata and a calibrated task-based model, we show these narratives reflect different channels through which AI affects the economy. Unlike previous waves of automation that increased both wage and wealth inequality, AI could reduce wage inequality through the displacement of high-income workers. However, two factors may counter this effect: these workers' tasks appear highly complementary with AI, potentially increasing their productivity, and they are better positioned to benefit from higher capital returns. When firms can choose how much AI to adopt, the wealth inequality effect is particularly pronounced, as the potential cost savings from automating high-wage tasks drive significantly higher adoption rates. Models that ignore this adoption decision risk understating the trade-off policymakers face between inequality and efficiency.

Keywords: Artificial intelligence, Employment, Inequality
JEL Codes: J21, O33, D31

*We would like to thank Nick Bloom, Pete Klenow, Chad Jones, Adrien Auclert, Chris Tonetti, Lukasz Drozd, Pierre-Olivier Gourinchas, Florence Jaumotte, and Giovanni Melina for helpful comments. We would also like to thank participants of the IMF Jobs and Growth seminar, the Stanford University macro workshop, as well as the 2024 Wharton AI and the Future of Work Conference among others, for valuable feedback.

[†]Stanford University, Department of Economics
[‡]International Monetary Fund

1 Introduction

Artificial intelligence (AI) is likely to profoundly affect many aspects of the economy, but its implications for inequality are highly debated. One common narrative is concerned with AI's potential to bring about widespread economic disruptions and magnify existing disparities. Other commentators have highlighted that, in contrast to previous episodes of automation, AI technologies are likely to predominantly disrupt high-income "white-collar" jobs, which could in fact lead to lower inequality. Finally, so-called "techno-optimists" focus on AI's potential to generate broad-based economic gains.[1] How can these competing narratives be reconciled?

In this paper, we ask: what will be the net effect of AI on inequality, and how might this compare to previous waves of automation? We investigate both theoretically and empirically how AI affects inequality through different channels – principally its impact on wages through labor displacement, and on wealth through capital returns. We find that while AI may reduce *wage* inequality by displacing high-income workers, it is likely to substantially increase *wealth* inequality as these same workers benefit from higher returns on their capital holdings. Moreover, our analysis suggests that high-income workers' occupations are particularly complementary with AI, meaning they may see productivity gains that could more than compensate the loss in wages due to task displacement.

It is worth clarifying how we conceptualize the distinction between AI and previous episodes of routine-biased automation. Throughout much of our analysis, we do not assume these are fundamentally different technologies in their economic mechanisms—both can be modeled as increasing the share of tasks performable by capital within occupations. This approach allows us to demonstrate that even if AI operates through similar economic channels

[1]For instance, two examples of articles sounding alarm bells on the potential increase in inequality appeared in the MIT Technology Review on April 2022 ("How to solve AI's inequality problem") and in Scientific American in August 2023 ("Unregulated AI Will Worsen Inequality, Warns Nobel-Winning Economist Joseph Stiglitz"). Meanwhile, in January 2023 The Atlantic published an article titled "How ChatGPT Will Destabilize White-Collar Work." and in July 2023 CNBC ran the headline "A.I. is on a collision course with white-collar, high-paid jobs — and with unknown impact". Finally, in September 2023 Forbes published an article titled "How To Make A Real Case For Technological Optimism".

as previous automation technologies, the difference in which occupations are exposed can still generate substantially different inequality outcomes. A key novelty of AI may not lie in how it functions economically, but rather in which segments of the labor market it primarily affects. That said, we do explore one potential technological distinction in our complementarity analysis, where we allow for AI to be more complementary with certain workers rather than purely substitutive, a characteristic that may differentiate it from earlier automation waves which primarily replaced routine tasks.

Using household microdata from the UK, we document that high-income workers are much more likely to work in occupations exposed to new AI technologies. We show that while roughly 60 percent of workers at the 90th income percentile are in an occupation where a large share of tasks can be performed by AI, at the 10th percentile only 15 percent of workers are in this situation. At the same time, workers in the upper-income decile also have the lowest share of their total income from wages, the largest wealth holdings, and the largest share of their wealth in risky but high-return assets such as firm equity. As such, high-income workers are not only better placed to insure themselves against potential adverse labor market impacts from AI, but also stand to benefit most from a possible AI-driven rise in capital returns. The picture is further nuanced by our finding that these workers are highly complementary with AI. While automation primarily substituted for low-wage routine tasks, AI may augment rather than replace many high-wage cognitive tasks, potentially increasing these workers' productivity.

To study the *a priori* ambiguous implications of AI adoption for inequality, and compare them to previous episodes of routine-biased automation, we thus consider a model that captures three key channels: i) the negative impact on wages due to heterogeneous job displacement across the income distribution, ii) the positive impact on capital income due to higher capital returns, and iii) the positive impact on wages due to aggregate productivity gains. This framework provides insights into each individual channel, their quantitative strength, as well as their joint effect.

Our starting point is the model of Moll et al. (2022), which we adapt to study AI diffusion instead of routine-biased automation, calibrated to the UK. [2] In this general-equilibrium model of heterogeneous households with a task-based production function, technological innovation acts as a shock to the share of tasks that can be performed by capital within each occupation. Households are heterogeneous in their income, wealth, and occupational exposure to new technologies. The higher capital share induced by a new technology thus impacts wage inequality by differentially displacing workers from tasks across the income distribution according to their exposure. Concurrently, it also affects wealth inequality via a capital income channel, because, as capital supply elasticity is limited, a rise in the capital share is accompanied by higher capital returns.

We use this model both to examine what impact automation had on wage and wealth inequality in the UK under this framework and also what effect AI adoption might be expected to have for technology adoption of a similar order as observed under previous episodes of automation. Under the baseline assumption that occupational exposure to AI is directly associated with task displacement, the model predicts a decrease in the Gini coefficient for wage inequality of 1.73 p.p.. This is driven both by an increase in wages for low-income workers due to aggregate productivity gains, and a decrease in wages for high-income workers due to task displacement. Wealth inequality, however, is predicted to widen, with the wealth Gini rising 7.18 p.p. For comparison, the model calibrated to routine-biased automation produces a substantial increase in both wage and wealth inequality, with the Gini rising by 2.05 p.p. and 6.89 p.p., respectively.

We can additionally allow for complementarity between labor and AI, which would compensate some exposed occupations in the form of higher productivity gains. When we do so, the model predicts a milder fall in wages for high-income workers, but also smaller wage gains (or even declining wages) for low-income workers, limiting the extent to which wage

[2] We choose the UK as our setting as it is an advanced economy with one of the highest degrees of exposure to AI (Pizzinelli et al., 2023), and has rich household data available, including detailed wealth holdings.

inequality may shrink (if decreasing at all). Moreover, the increase in aggregate productivity due to AI would partially mitigate the negative effect of task displacement on wages, thus increasing the share of workers who are unconditionally better off from AI diffusion, and further raising capital returns. However, this aggregate improvement does not fundamentally alter the inequality effects of AI adoption, with wealth inequality still increasing markedly compared to wage inequality.

Our findings naturally raise the question of whether policymakers should attempt to offset the likely negative impacts of AI on wealth inequality using redistributive policies such as capital taxes. Our baseline model (and the broader literature it follows) cannot speak to such questions because taking the AI-induced increase in the capital share as exogenous does not allow for such policies to impact AI adoption. In reality, however, policies such as a capital tax would pose an important trade-off for policy makers: they provide an opportunity to reduce the potentially large inequality impact of AI, but they may also hamper productivity growth by creating misallocation and limiting AI adoption.

To speak to this trade-off, we develop an extended model with an endogenous capital share and endogenous AI adoption decisions by firms. This model allows us to analyze how firms' endogenous adoption decisions interact with worker characteristics to amplify the differential impact of AI relative to previous waves of automation. We find that the potential cost savings from automating high-wage tasks lead to significantly higher AI adoption compared to our baseline model, resulting in larger productivity gains but also more pronounced effects on wealth inequality. This amplified adoption intensifies the trade-off policymakers face between reducing inequality and promoting growth. We find that while a 15 percent capital tax could substantially reduce post-transfer inequality – lowering the wage Gini by 3.4 p.p. and the wealth Gini by 3.7 p.p. – it would come at significant economic cost, reducing output by 26.9 percent and average wages by 11.8 percent. These costs are nearly twice as large as they would have been for automation, suggesting policymakers may need to carefully consider the growth implications of policies aimed at addressing the distributional concerns that arise from these new technologies.

There is a large literature on automation and its impact on labor markets. Zeira (1998) was the first paper to model automation as a shock to workers via displacement in a task-based production framework. Following many others in the literature (Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018; Moll et al., 2022), in this paper, we similarly model technological change via a task-based framework. One of our contributions relative to the existing literature is to consider the implications of this canonical framework for AI, given the different incidence of exposure to previous episodes of routine-biased automation, as documented by Felten et al. (2021). Relative to this and other recent papers in this literature focused on exposure to AI (Acemoglu et al., 2022; Brynjolfsson et al., 2023) we also explicitly account for the possibility that AI may serve as a complement to some workers, rather than just a substitute. Building on the work of Pizzinelli et al. (2023) we show that this has stark implications for the potential inequality implications of these new technologies. Our findings complement recent policy work highlighting both AI's potential to boost productivity and its distributional implications, particularly the risks of rising inequality if complementarity with high-income workers is strong (Cazzaniga et al., 2024), as well as the need for fiscal policy to help steer the technology's development and cushion its negative effects on workers (Brollo et al., 2024).

Our paper also contributes to the broader literature on the drivers of technology adoption. While a large number of studies have examined why firms may adopt a new technology, how technologies spread, and the role of labor markets in technology diffusion (Buera and Oberfield, 2020; Perla et al., 2021; Benhabib et al., 2021; Humlum, 2021), the existing literature on automation largely takes the adoption decision as exogenous. In existing models of technological change in a task-based framework, automation is typically represented by a change in the share of tasks capital can perform, where capital will always be chosen to perform a task if available as it is assumed to be lower cost. In contrast, we show in our extended model how we can expand this framework to allow for endogenous adoption decisions, building on the theory of Drozd et al. (2022). This model highlights how the unique characteristics of AI may lead to amplified adoption and more pronounced economic impacts compared to previous waves of automation. We argue that such a model is essential for thinking about how policymakers may respond to the advent of new technologies, and what optimal redistributive policy may look like, given the potential trade-off for the policymaker between inequality objectives and the potential productivity gains of adoption.

Finally, related to our contributions above, we also add to existing work on the relationship between technological change and inequality (Autor et al. (2003); Autor et al. (2006); Korinek and Stiglitz (2017); Hémous and Olsen (2022); Moll et al. (2022); Acemoglu and Restrepo (2022)). First, our findings on the differential inequality impacts of AI relative to previous episodes of automation, and the importance of complementarity for the net effect are novel. Moreover, a conceptual contribution of our model extension is to highlight the interaction between firms' adoption decisions and redistributive policies, producing a potential trade-off between equity and aggregate welfare and productivity.

The paper proceeds as follows. In the next section, we describe how we define and measure exposure to a technology (both automation and AI), as well as potential AI complementarity, and we discuss the UK household microdata. Section 3 contains our main empirical results. Section lays out our baseline model, calibration and results, as well as exploring the role of complementarity. We then show how this can be extended in Section 5 to endogenize the capital share and technology adoption decision, and show what this extended model would predict for AI adoption. In Section 6, we then use this model with endogenous technology adoption to examine the trade-off between equity and efficiency of redistributive policies. Finally, Section 7 concludes.

2 Data

2.1 Measuring Technology Exposure

To inform the empirical analysis and the calibration of the model we use measures of exposure to AI and routine-biased automation developed by previous studies. These measures provide a taxonomy to identify occupations that have a greater *ex ante* degree of exposure to these technological changes based on their characteristics, such as the key tasks to be carried out in a job and the abilities needed to perform them successfully. Both measures were developed using O*NET, a publicly available repository codifying occupations in the US.

To study exposure to automation we construct a routine task index (RTI) similar to Autor and Dorn (2013), and based on the definition of routine jobs in Acemoglu and Autor (2011), using the O*NET classification of routine tasks. This index measures the exposure to automation by occupations' relative intensity of routine tasks compared to manual and cognitive ones. In this context, the process of routine-biased automation is assumed to lead to a greater displacement of labor in occupations with a higher share of routine tasks, as evidenced by a large literature for both the US and the UK, starting with Autor et al. (2003) and Goos and Manning (2007). To study exposure to AI, we consider the AI Occupational Exposure (AIOE) index proposed by Felten et al. (2021). This index considers the overlap between AI applications in several fields and the human abilities needed to perform a given occupation, thus appraising the degree to which AI can replicate the skills essential to each job.

Unlike the RTI, the AIOE measure remains agnostic regarding whether AI exposure necessarily entails substitution of human labor within each job or could function as a support technology for workers. In our baseline scenario, we interpret AI exposure as reflecting solely the potential for worker substitution. However, in an alternative scenario, we make a distinction between labor-substituting and labor-augmenting exposure. To this end, we introduce the concept of potential complementarity, following Pizzinelli et al. (2023). This measure considers a set of social and environmental contexts related to each job reported in O*NET, which may mitigate the likelihood of AI replacing labor. We use this complementarity index to more explicitly differentiate between the risk of labor substitution from the potential increases in productivity related to AI exposure.

2.2 Household Data

Our empirical analysis, which motivates the theoretical model and informs its calibration, is conducted using the Wealth and Assets Survey (WAS) for the UK. This dataset is a large nationally-representative survey of the financial holdings and sources of income of households in the UK.

The survey contains both person- and household-level information relevant for our analysis. From the person-level dataset we use variables related to individuals' employment status, occupation, wages and hours worked per week. We also examine other sources of income, including government transfers, public and private pensions, property income, and investment income. Through this information, we construct the distribution of workers' total income and examine its composition from different sources across different income quantiles. The survey also contains detailed information on the asset holdings of each individual or of the household (if jointly owned). We use this information to study the composition of households' balance sheets across the income distribution.

For the analysis, we use the 2016, 2018, and 2020 waves of the survey (covering April 2016-March 2020), restricting the sample to working-age individuals who are employed.[3] Our final sample is comprised of 28,588 workers across 14,780 households. All nominal values in British Pounds are deflated to 2019 prices.

[3]In our analysis we focus on employees and exclude self-employed individuals and entrepreneurs. This allows us to avoid measurement issues around self-employed incomes, and also allows us to focus on 'workers' who risk being displaced by AI and who's main income source is their labor income, as opposed to entrepreneurs who receive a mix of labor and capital income. By only using data up to March 2020, we avoid the temporary impact of COVID on sectoral composition, working hours, income and savings, as some occupations and sectors experienced significantly more disruption than others.

3 Empirical Analysis

3.1 Exposure Across the Income Distribution

From the WAS, we use information on workers' occupations to compare automation exposure to AI exposure across the income distribution. To this end, we convert the continuous exposure measures described in Section 2 into binary variables following the approach of Moll et al. (2022). We first sort occupations by their routine task intensity, and define as exposed to automation those occupations with the highest values of the index that comprise 30 percent of total hours worked in the UK across 2016-2020. We repeat the same process to classify the occupations exposed to AI. We use 30 percent as a threshold for comparability with Moll et al. (2022), and keep the same threshold for both automation and AI for comparability across the two technologies. This approach naturally implies that at the aggregate level, economy-wide exposure to automation and AI are assumed to be of equal magnitude. We then calculate the share of workers in exposed occupations for each percentile of the distribution of total income (including both labor and non-labor sources). [4]

Figure 1 shows how exposure to automation and AI varies by income percentiles. For automation, consistent with the general narrative that low-income workers were worst affected, we can see clearly that exposure is broadly decreasing in income. Nearly 50 percent of workers in the lowest income percentiles were exposed to automation, compared to fewer than 20 percent of high-income workers. In stark contrast, it is largely high-income workers who appear to be most exposed to the new AI technologies. Fewer than 20 percent of workers in the bottom income decile are exposed to AI, compared to over 60 percent in the top income decile.

[4] This approach also treats exposure as homogeneous within occupations, while in reality, there is likely significant heterogeneity in how AI affects different tasks and workers within the same occupation. Even in highly exposed occupations, certain high-income tasks may be highly substitutable while others remain complementary with AI. Our occupation-level analysis should therefore be viewed as a first-order approximation of the distributional impacts of AI adoption, potentially overlooking some of the more nuanced within-occupation effects.

Figure 1: Exposure to Automation and AI by Income Percentile

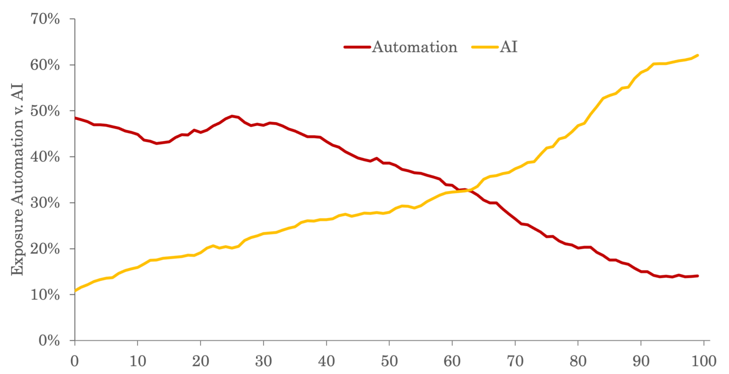

Notes: Each worker is classified as 'exposed' if they work in the top 30 percent of occupations (weighted by hours worked) of share of tasks that are routine (for automation) or share of tasks that can be performed by AI (for AI). The average share of exposed workers is then calculated for each income percentile (as measured by total income).

While Figure 1 shows worker-level exposure to these technologies, the WAS also allows us to examine the correlation of exposure between members of the same household and the overall household-level exposure. If, for example, the within-household exposure correlation is low, then these stark income trends in exposure may have milder implications for household-level income inequality compared to individual-level disparities. However, consistent with assortative matching, Appendix Figure B.2 shows that there is a high degree of correlation in exposure within households, and therefore very similar patterns in individual and household exposure. As a result, we focus on worker-level characteristics and exposure for the remainder of our analysis.

We can use the complementarity index developed by Pizzinelli et al. (2023) to consider how the extent to which workers may be complementary with new AI technologies varies across the income distribution. Doing so nuances the implications of the potential for displacement across the income distribution. This is because, as shown in Figure 2, although the highest-income workers are the most exposed to AI, they are also much more likely to be complementary with these new technologies. Complementarity monotonically increases in income up until the 80th percentile, after which it levels off. This suggests any analysis comparing the inequality impact of AI versus automation should consider the possible mitigation of job displacement and even the potential benefits of increased productivity implied by this dispersion in complementarity.

Figure 2: AI Complementarity Index by Income Percentile

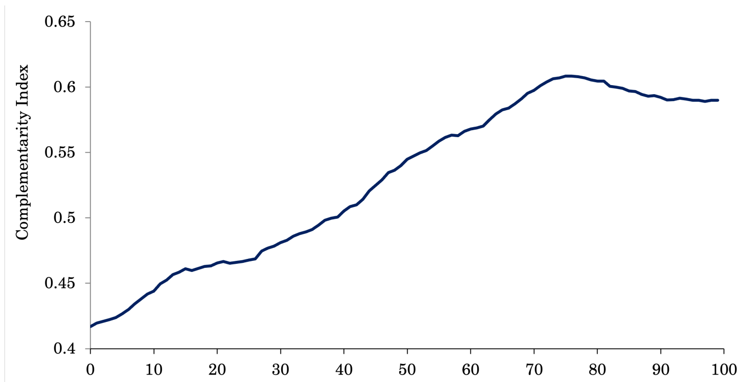

Notes: We calculate the complementarity of each occupation following the measure developed in Pizzinelli et al. (2023). The average complementarity of workers' occupations is then calculated for each income percentile (as measured by total income).

3.2 Income and Wealth

The WAS also allows us to consider workers' other income sources across the earnings distribution. This is informative of their ability to self-insure from any labor shocks resulting from AI exposure, as well as the potential to benefit from the impact these technologies may have on assets related to firms' profits, such as equities.

Figure 3 highlights that the highest-income workers, who face the highest exposure to AI, may have more options to mitigate any potential negative employment consequences, as they have the largest share of income from non-labor sources. In addition to their wage income, they also receive substantial "capital income" such as rent on properties and investment returns, and pension income (which will in large part reflect returns on assets). Moreover, the investment income captured in the WAS only reflects realized investment income such as dividends, but not unrealized capital gains for increases in asset prices, meaning this income source is likely to be significantly understated.

To speak to that, we directly examine how wealth holdings and portfolio shares vary

Figure 3: Share of Total Income by Income Percentile

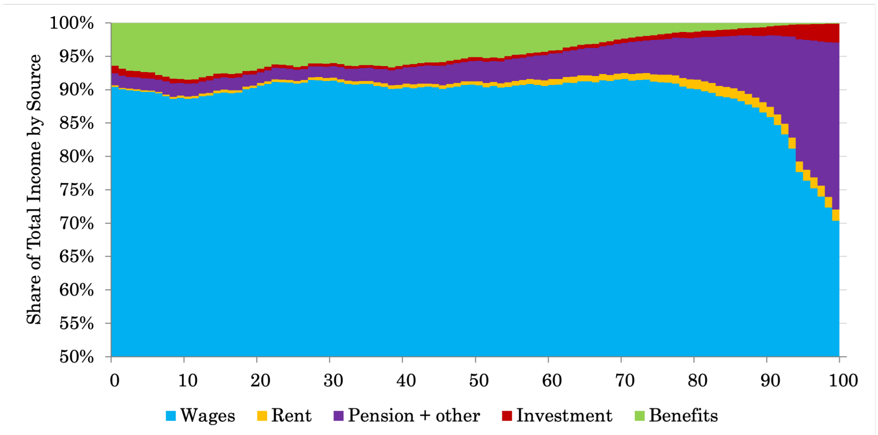

Notes: Total income reflects net individual income, plus benefits, from the listed sources. It also captures household income divided equally among all adults in the household (for example from rent on shared property).

across the income distribution.[5] Figure 4a shows that wealth is highly concentrated among high-income workers. Furthermore, Figure 4b shows the share of workers' wealth that is invested in "risky assets", those more likely to see an AI-driven rise in value, is much higher for high-earners than for low-income workers.[6]

Overall, we find that, in contrast to automation, high-income workers are much more exposed to AI than low-income workers. This may suggest a potential for AI to reduce wage inequality via labor displacement. However, this prediction may change when accounting for the fact that high-earners are also more likely to benefit from AI, both due to potential complementarity and through higher returns to their capital holdings. Therefore, the net effect of AI on inequality therefore depends on a number of competing channels, motivating the need for a structural framework.

[5]To measure workers' wealth, we take individual wealth and supplement it with jointly-owned household wealth (such as the value of the main residence), divided equally among adults in the household, where appropriate.

[6]It is important to note that we carefully distinguish between different types of pension income in this analysis. The UK has both defined benefit (DB) and defined contribution (DC) pension schemes. For our analysis of portfolio choice and when calibrating the model, we classify only defined contribution pensions as risky assets that would benefit from higher equity returns potentially driven by AI adoption. Defined benefit pensions, which guarantee fixed payments regardless of asset performance, are treated as safe assets (similar to government bonds) since households with these pensions would not directly benefit from higher capital returns.

Figure 4: Asset Portfolios by Income Percentile

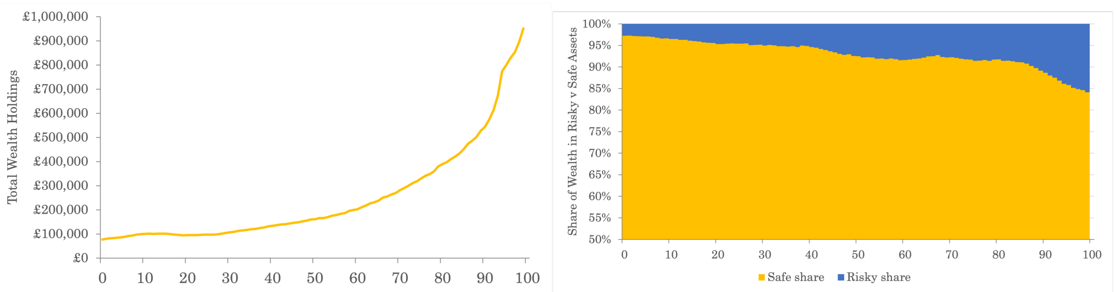

(a) Total Wealth by Income Percentile (b) Risky vs. Safe Asset Shares

Notes: "Safe assets" include property wealth, physical wealth, defined-benefit pensions, cash and government bonds. "Risky assets" include defined-contribution pensions, equities, and corporate bonds.

4 Structural Analysis

To assess the net impact of AI on inequality, and how it compares to automation, requires bringing together three channels into a quantitative framework: the negative impact of task displacement on real labor income; the positive impact of productivity gains on real labor income; and the positive impact of higher returns on capital income. Given the findings of Section 3, it is likely that the strength of each channel will differ across the income distribution, and differently for automation and AI, and as such our structural framework must be able to capture this heterogeneity. To analyze these competing channels quantitatively, we develop our structural analysis in two stages. We begin with a baseline model that, following the existing literature, takes technology adoption and the resulting changes in capital shares as exogenous. This allows us to isolate the direct effects of AI exposure on inequality through labor displacement and capital returns. We then extend the framework in Section 5 to endogenize firms' adoption decisions and capital use, showing how the cost savings from automating high-wage tasks could amplify both the productivity gains and distributional impacts of AI.

4.1 Baseline Model

We take as our starting point the model of Moll et al. (2022). As the model is described in great detail in the original paper, below we focus on the parts that are most important for the channels we study and provide a summary description of the other components.

In this general equilibrium model of heterogeneous households with a task-based production function, technological innovation acts as a shock to the share of tasks that can be performed by capital within each occupation. Specifically, households consume a final good Y, which is produced from the intermediate outputs of many sectors, indexed by z, via Cobb-Douglas aggregation:

$$Y = A \prod_z Y_z^{\eta_z} \tag{1}$$

where η_z gives the share of each intermediate good in the production of the final good, and $\sum_z \eta_z = 1$.

Each intermediate good is itself the Cobb-Douglas aggregate of a continuum of tasks u:

$$\ln Y_z = \int_0^1 \ln \mathcal{Y}_z(u) du \tag{2}$$

where we assume that a share, α_z, of tasks can be performed by capital within each sector z, and ψ_z represents the efficiency of labor at performing a given task z:

$$\mathcal{Y}_z(q) = \begin{cases} \psi_z l_z(u) + k_z(u) & \text{if } u \in [0, \alpha_z] \\ \psi_z l_z(u) & \text{if } u \in (\alpha_z, 1]. \end{cases} \tag{3}$$

In the baseline model, if capital is able to perform a task, firms will choose to use it as it is assumed to be cheaper than labor, with its relative cost determined by the parameter ψ_z, which we discuss further below. This means that the capital share within a sector is simply given by α_z.

In this framework, we can think of technological change as an increase in the share of tasks that can be performed by capital. By mapping each percentile of our worker income distribution into a sector z, we capture heterogeneity in the exposure of workers to different technologies as heterogeneity in the sector-level changes in the α_z's resulting from the adoption of the technology. We discuss the calibration of α_z in Section 4.2.

In addition to capturing the heterogeneity in technological exposure documented in Section 3, another advantage of this framework is that it captures heterogeneity in the potential benefits households may receive from higher capital income. This is because households, in addition to being heterogeneous in their occupation (and therefore labor income), also differ in their wealth and portfolio holdings (and therefore capital income).

Moreover, this model incorporates a key empirical feature: the capital supply elasticity with respect to the return on capital is finite rather than infinite. This means that when technological innovation increases the demand for capital, the capital supply does not fully adjust even in the long run, resulting in a higher equilibrium return on capital. To generate this realistic feature in a tractable way, the model (following (Moll et al., 2022)) assumes households face some probability of a "wealth dissipation shock" where their stock of assets fully disappears. Although this is an unrealistic assumption in itself, it serves as a tractable proxy for various factors that prevent perfect capital accumulation in response to higher returns, such as technological obsolescence, imperfect intergenerational transfers, or demographic changes. The advantages of this approach are that it yields a non-degenerate wealth distribution and, importantly for our analysis, it delivers a non-degenerate wealth distribution, and importantly for us, an increase in the capital share due to technological innovation will increase the return on capital, capturing our third channel of impact.

Households have Epstein-Zin preferences over their consumption:

$$v_0 = E_0 \int_0^\infty \frac{\rho(1-\gamma)v}{1-\sigma} \left(\left(\frac{c}{((1-\gamma)v)^{1/(1-\gamma)}} \right)^{1-\sigma} - 1 \right) dt$$

$$da_{z,t} + db_{z,t} = (r_K a_{z,t} + r_B b_{z,t} + w_z - c_{z,t})dt + a_{z,t} v dW_t$$

where $\rho = \varrho + p$ represents households' discount rate, which is a function of their impatience (ϱ) and the probability that they are hit with a wealth dissipation shock (p). The parameter γ captures their risk aversion, and σ is the inverse intertemporal elasticity of substitution. Finally, r_K is the return on capital and r_B is the return on bonds, where only a fraction χ of households can invest in capital.

Note that households inelastically supply labor to their assigned sector z, an important simplification for tractability. This means that our model will feature no unemployment; workers will be displaced from tasks but not from jobs.[7] However, it also means that in response to a shock, workers cannot, for example, switch occupations. Our results should therefore be considered as the impact of a shock in the face of significant switching costs, or in the absence of any policies to promote retraining. Of course, such policies may provide substantial welfare benefits to the extent to which they are able to help workers avoid a large degree of task displacement, and workers may also differ in their ability to retrain and take advantage of new technologies. Although beyond the scope of this paper, this is something discussed in more detail in Cazzaniga et al. (2024).

4.2 Calibration of Technology Adoption

The most important parameters for the analysis are those related to the change in the sector-level capital shares, captured by α_z. These will determine how much displacement is induced by automation and AI, respectively, in each sector and need to reflect the exposure patterns documented in Section 3.

To calibrate these, we follow Moll et al. (2022) in using a shift-share approach, based on the change in the aggregate capital share. For example, for automation, we relate the

[7]Note that task displacement means that a worker will continue in their existing role, but demand for their labor will decline, so all else equal their labor income will fall.

evolution of the aggregate and sectoral capital shares as follows:

$$\frac{1}{1-\alpha_{z,2014}} - \frac{1}{1-\alpha_{z,1980}} = \omega_z^R \left(\frac{1}{1-\alpha_{2014}} - \frac{1}{1-\alpha_{1980}} \right).$$

This apportions the change in the aggregate capital share ($\left(\frac{1}{1-\alpha_{2014}} - \frac{1}{1-\alpha_{1980}}\right)$) into changes in sectoral capital shares based on the routine exposure of each sector, represented by ω_z^R. Consistent with Moll et al. (2022), the time interval chosen, 1980-2014, corresponds to the decades over which automation adoption was most pronounced in industrialized economies.

We follow the same strategy for AI, using the exposure index of Felten et al. (2021): ω_z^{AIOE}. However, since we are using the model as a forward-looking tool, we do not know what the aggregate change in the capital share induced by AI will ultimately be. For our baseline we therefore assume that the change in the aggregate capital share in response to AI will be of the same magnitude as that observed for automation; that is, we assume the change in the aggregate capital share over 2014-2048 will match the change in the aggregate capital share from 1980-2014. This allows for an immediate comparison of the inequality impacts of automation versus AI by holding the aggregate magnitude of the shock fixed. We also perform sensitivity analysis in Appendix B showing how different changes in the aggregate capital share affect our findings.

The other important set of parameters for the analysis of AI adoption are those that govern sectoral complementarity with AI. In the alternative scenario, we model complementarity as a change in the sector's weight in the Cobb-Douglas production function η_z in response to technology adoption, consistent with the idea that the value-added of the sectors with greater complementarity increases. If η_z increases for a sector, even if workers employed in sector z perform fewer tasks, the tasks they do perform are now worth relatively more. To see this more intuitively, consider the equation for wages within a sector:

$$w_z \; = \; (1-\alpha_z)\, \frac{\eta_z}{l_z}\, Y(K) \tag{4}$$

Technological innovation can exert a downward pressure on wages within a sector via a higher α_z, but also potentially push them up via a higher η_z. We calibrate the initial η_z to match the wage distribution in 2014. For automation and our baseline AI scenario we keep these parameters fixed following technology adoption. Meanwhile, in the scenarios where we allow for complementarity we calibrate an increase in η_z as: $\eta_{z,T} = \eta_{z,0}(1 + \lambda)$, where λ captures the relative complementarity of different occupations. Note that because η_z are sectoral weights, they still must sum to one, so after recalibrating them we renormalize them. This implicitly means that complementarity shifts the relative weights in production such that some sectors will experience higher income (and therefore higher wages) at the expense of other sectors. While this reallocation is likely to be an important aspect of complementarity, we also consider a second complementarity scenario below that allows for complementarity to additionally increase aggregate productivity in a way that makes all sectors better off. In this scenario, in addition to changing the sectoral weights, complementarity also increases ψ_z such that: $\psi_{z,T} = \frac{w_z}{R}/1.3 * (1 + \theta_z - \theta_{min})$.

We can then calibrate the remaining model parameters for the UK using estimates from the literature, as well as moments from the WAS and UK aggregate data. When calibrating wealth holdings, we account for the different nature of pension assets. Only defined contribution pension wealth is included in our measure of risky assets, as these directly expose households to changes in capital returns. Defined benefit pension entitlements are treated as safe assets, reflecting their fixed payment structure that is largely insulated from current market returns. This gives us the parameters as described in Table 1.

Table 1: Calibration of Remaining Parameters

Parameter:	**Description**	**Value**	**Source**
σ	Inverse IES	2	Standard
γ	Risk aversion	2	Standard
p	Dissipation rate	4.5%	Moll et al.
ρ	Discount rate	1%	Moll et al. (target r=6.5%)
ξ	Share investors	10%	Target $\kappa_{2014} = 1.35$ and $\frac{1}{\zeta_{2014}} = 0.46$
ν	Capital risk	6%	Target $\kappa_{2014} = 1.35$ and $\frac{1}{\zeta_{2014}} = 0.46$
g	TFP growth in ψ_z	1.5%	Standard
δ	Depreciation rate	5%	Standard
A	Productivity term	0.143	Y/L in 2014
$\eta_{z,0}$	Sector shares	-	Match income distribution in 2014
$\psi_{z,0}$	Relative productivity of labor	-	Moll et al. (cost savings of 30%)

4.3 Baseline Results

Using the model, we can now do a steady state comparison of wages and total income across the income distribution before and after the adoption of the new technologies. Figure 5 shows how wages and total income changed between 1980 and 2014 in the case of automation. We can see in Figure 5a that task displacement fell predominantly on low-income households, suppressing their wages.[8] At the same time, high-income households also saw a substantial increase in their capital income (shown in red in Figure 5b, with labor income in blue), particularly the top 1 percent. As a result, automation substantially increased both wage and total income inequality in the UK. We calculate that the wage Gini rose 2.05 p.p., and the wealth Gini rose by 6.89 p.p. due to automation.

In contrast, for our baseline AI scenario (where we do not allow for complementarity), high-income workers are the most likely to see their wages fall in response to AI displacement, whereas low-income workers see wage gains due to higher aggregate productivity (Figure 6a). As a result, wage inequality falls in this scenario, with the wage Gini falling 1.73 p.p. over

[8]In Appendix B we compare how the distribution of wages predicted by the model compares to the change in wages actually observed in the UK between 1980-2014.

Figure 5: Impact of Automation on Wages and Total Income, 1980-2014

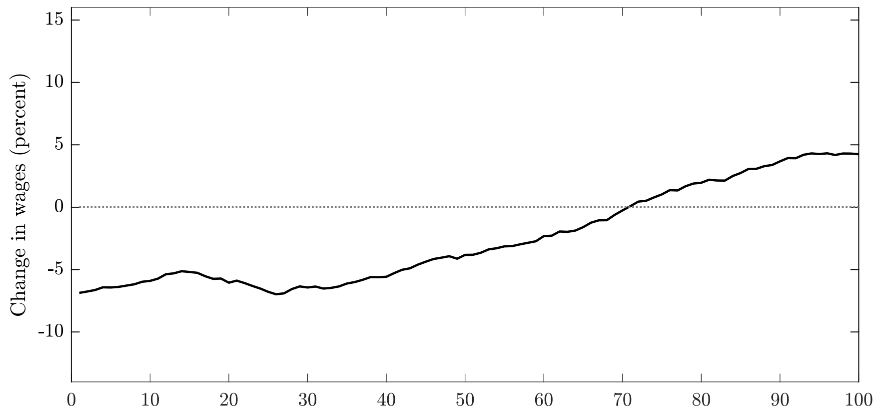

(a) Change in Wages by Income Percentile

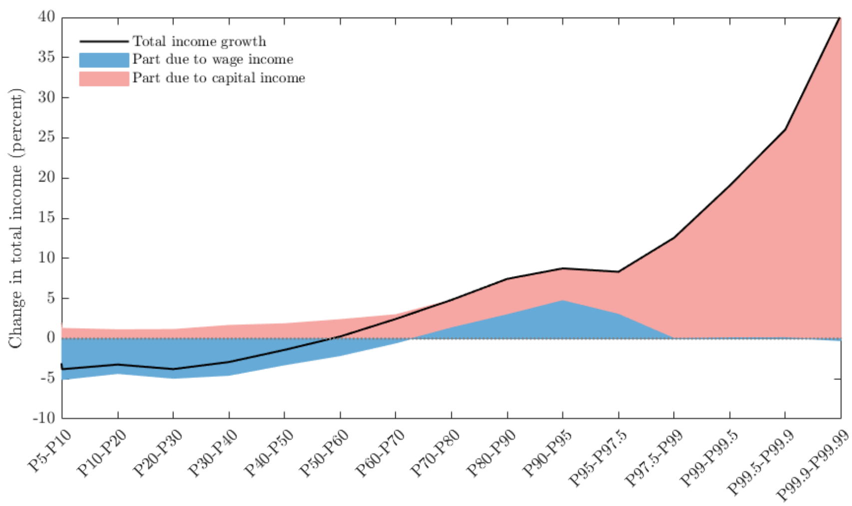

(b) Change in Total Income by Income Percentile

Notes: Figure (a) shows the percent change in wages between 1980 and 2014 by total income percentile. Figure (b) shows the percent change in total income between 1980 and 2014 by total income percentile, decomposed into the change due to labor income in blue, and the change due to capital income in red. Note that the scale for Figure (b) is nonlinear, with the top-tail of the income distribution expanded.

the scenario. However, this fall in wages is more than offset for the highest-income workers by higher capital income, meaning their total income is still expected to rise significantly (Figure 6b). As a result, wealth inequality is still expected to increase, with the wealth Gini rising 7.18 p.p. over the scenario.

Figure 6: Predicted Impact of AI on Wages and Total Income, 2014-2048

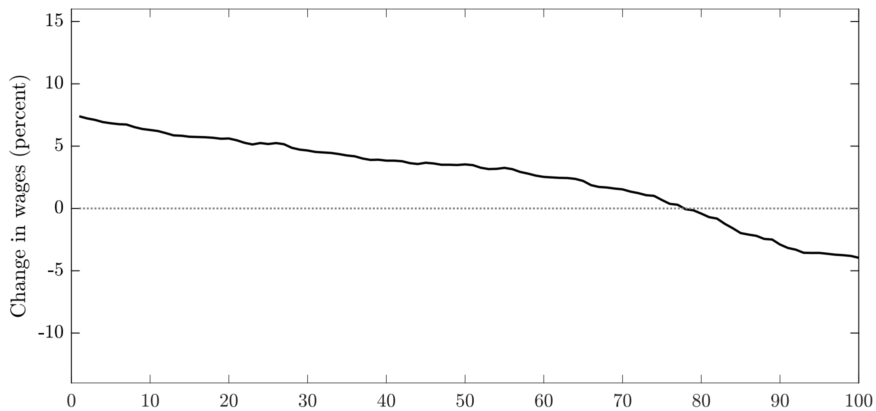

(a) Change in Wages by Income Percentile

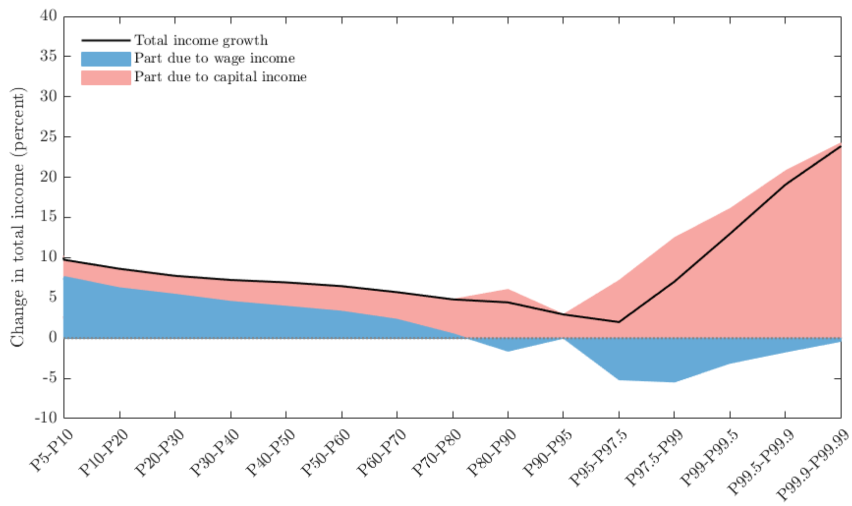

(b) Change in Total Income by Income Percentile

Notes: Figure (a) shows the percent change in wages between 2014 and 2048 by total income percentile. Figure (b) shows the percent change in total income between 2014 and 2048 by total income percentile, decomposed into the change due to labor income in blue, and the change due to capital income in red. Note that the scale for Figure (b) is nonlinear, with the top-tail of the income distribution expanded.

4.4 Accounting for Complementarity

When allowing for potential complementary with AI, with sectoral heterogeneity modeled as described above, the distinctions between automation and AI become more nuanced. Given that middle- and high-income workers have the highest expected complementarity, even though these workers see downward wage pressure from displacement, they also benefit from an increase in their relative productivity, which results in wage gains. At the same time, the low complementarity of low-income workers means a lower share of productivity gains accrue to them, thereby lowering their wages. As a result, in this scenario the reduction in wage inequality is much more muted, with the wage Gini only expected to fall by 0.22pp, while wealth inequality is expected to increase to a similar degree (Table 2).

Figure 7: Predicted Impact of AI on Wages and Total Income, 2014-2048 Allowing for Complementarity

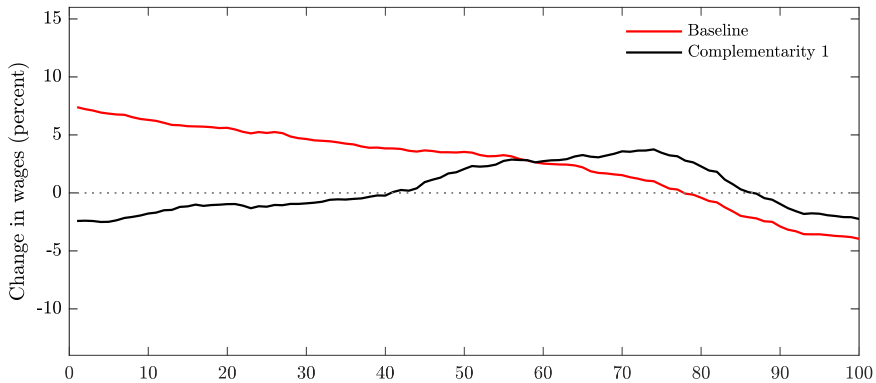

(a) Change in Wages by Income Percentile

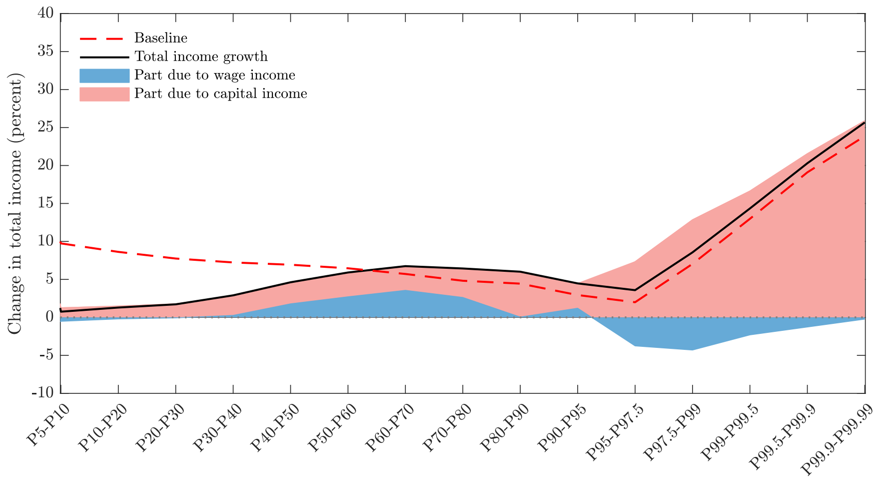

(b) Change in Total Income by Income Percentile

Notes: Figure (a) shows the percent change in wages between 2014 and 2048 by total income percentile for the baseline AI scenario in red, relative to the scenario allowing for complementarity in black. Figure (b) shows the percent change in total income between 2014 and 2048 by total income percentile in the baseline AI scenario in red, relative to the scenario allowing for complementarity in black, decomposed into the change due to labor income in blue, and the change due to capital income in red. Note that the scale for Figure (b) is nonlinear, with the top-tail of the income distribution expanded.

Finally, in addition to allowing for complementarity to affect the relative share of productivity gains that accrue to different sectors, we also allow AI to yield aggregate productivity gains via TFP growth. This assumption would reflect a view of AI as being more productivity-inducing than automation for the same change in the capital share. Moreover, combining this assumption with complementarity means that all sectors experience productivity improvements, but the gains are distributed unevenly, with the size of increase not necessarily proportional to sectors' initial productivity levels. To do so, in addition to recalibrating the sectoral weights as: $\eta_{z,T} = \eta_{z,0}(1+\theta_z)$, we also calibrate the average productivity gains conditional on adoption as $\psi_z = \frac{w_z}{R}/1.3 * (1+\theta_z - \theta_{min})$. Figure 8 and Table 2 suggest that under the addition of aggregate productivity gains all workers are better off from AI adoption, as higher TFP raises the base levels of all wages. However, in this scenario the inequality implications remain unchanged compared to the scenario with only complementarity and no productivity growth, as the differential effect on wages and total income across the distribution are the same.

While our model accounts for productivity gains through complementarity with existing tasks, it does not capture another important channel through which AI might be complementary with workers: the creation of entirely new products, services, and occupations. Historical technological disruptions, such as the internet revolution, led to entirely new job categories (e.g., social media manager, app developer) that could not have been predicted beforehand. As Acemoglu and Restrepo (2019) highlight, technological change often creates new tasks that can be performed by labor, potentially offsetting some of the displacement effects. Our estimates should therefore be interpreted as considering only the impacts on existing occupations, potentially understating the positive labor market effects of AI.

Table 2: Change in Inequality Metrics, Aggregate Productivity and Output

Scenario:	**Auto (baseline)**	**AI (baseline)**	**AI (comp. 1)**	**AI (comp. 2)**
Wage Gini	2.05 p.p.	-1.73 p.p.	-0.22 p.p.	-0.22 p.p.
Wealth Gini	6.89 p.p.	7.18 p.p.	7.16 p.p.	7.16 p.p.
Capital Share	5.5 p.p.	5.5 p.p.	5.5 p.p.	5.5 p.p.
Output	9.6%	10.6%	10.6%	12.2%
Mean Wages	0.2%	0.2%	0.2%	1.6%

Figure 8: Predicted Impact of AI on Wages and Total Income, 2014-2048
Allowing for Complementarity and Additional Productivity Gains

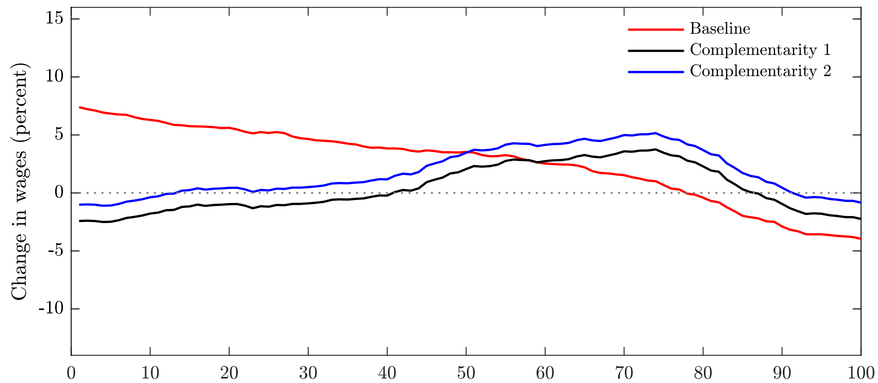

(a) Change in Wages by Income Percentile

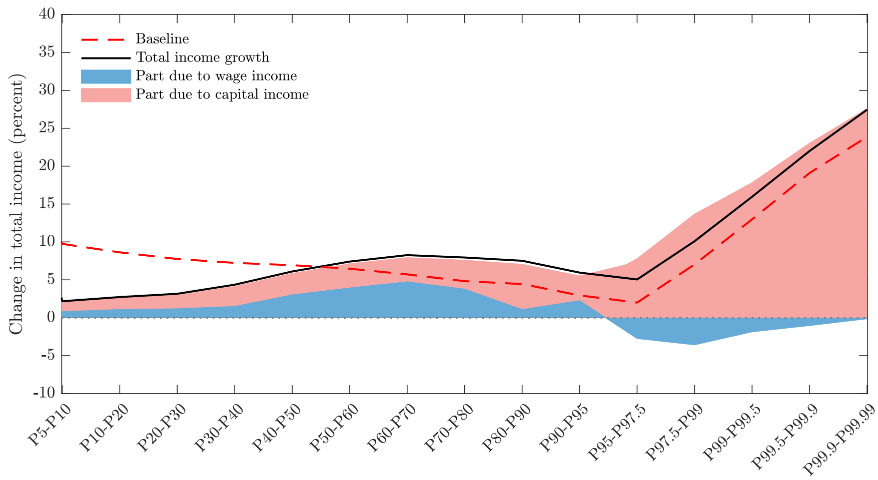

(b) Change in Total Income by Income Percentile

Notes: Figure (a) shows the percent change in wages between 2014 and 2048 by total income percentile for the baseline AI scenario in red, relative to the scenario allowing for complementarity in black, and the scenario allowing for complementarity with additional productivity gains in blue. Figure (b) shows the percent change in total income between 2014 and 2048 by total income percentile in the baseline AI scenario in red, relative to the scenario allowing for complementarity with additional productivity gains in black, decomposed into the change due to labor income in blue, and the change due to capital income in red. Note that the scale for Figure (b) is nonlinear, with the top-tail of the income distribution expanded.

5 Extended Model

In the baseline model, the adoption of new technologies, whether automation or AI, is taken as exogenous. Technological change is reflected in the share of tasks that could be performed by capital and, as capital is assumed to always be cheaper than labor, a firm would always choose to use it to perform a task if possible. In reality, however, the adoption of new technology is costly, and for some tasks the benefits of adoption are unlikely to outweigh the costs. This likely explains in large part why we see differential adoption of technologies across countries and sectors. For example, in countries with cheaper labor, limited access to new technologies, or where scant supporting infrastructure makes technology adoption costly, firms may optimally prefer to use labor to perform tasks even if they could in theory be automated.

To better understand this choice on the part of firms and the conditions that incentivize adoption, we need to move towards a framework in which firms' capital use and adoption decisions are endogenous and depend on these factors. Another advantage of extending the framework in this way is that it highlights an important trade-off policymakers face. A natural question given the findings from the baseline model is the extent to which policymakers can offset the impact of new technologies on inequality using redistributive measures. The baseline model (and others in the literature that take the capital share as exogenous) does not allow for meaningful consideration of such policies, as the exogeneity of the adoption decision means there is no scope for policies to distort firms' capital use or technology adoption decisions. In practice, policies such as a capital tax, while reducing inequality, also risk inducing misallocation and lower productivity gains from reduced adoption. By endogenizing the adoption decision, our extended framework below clearly highlights this trade-off.

5.1 Endogenizing the Initial Capital Share

For the firm's problem, we adapt our baseline model to endogenize the capital share prior to the adoption of new technologies. To do so, we maintain our assumptions that the economy produces a final good, which is itself a Cobb-Douglas aggregate of perfectly competitive intermediate goods, and households have Epstein-Zin preferences. The production of the intermediate goods occurs through the combination of a continuum of tasks, each individually performed by either capital or labor, and aggregated via a Cobb-Douglas function. Our key assumption is that, following Drozd et al. (2022), tasks vary in complexity with a Pareto distribution. Moreover, while the labor input required to complete a given task is independent of the complexity of the task, the amount of capital required for a task increases with the task's complexity. As a result, the optimal level of capital use for the production of an intermediate good will depend on the share of tasks for which capital is cheaper than labor.

Households consume the final good Y, which is produced from the intermediate outputs of each sector z via Cobb-Douglas aggregation:

$$Y = A \prod_z Y_z^{\eta_z} \tag{5}$$

where η_z gives the share of each intermediate good in the production of the final good, and $\sum_z \eta_z = 1$.

Each intermediate good is itself the Cobb-Douglas aggregate of a continuum of tasks, each with complexity q:

$$\ln Y_z = \int_0^{\infty} \ln \mathcal{Y}_z(q) dq \tag{6}$$

All tasks can be produced with either capital or labor. Moreover, while the labor input required to complete a given task is independent of the complexity of the task, the amount of capital required for a task is upward-sloping in complexity q:

$$k_z(q) = Z_z q^{\lambda} \tag{7}$$

$$l_z(q) = Z_z \tag{8}$$

As in the baseline model, the cost of a unit of capital input is the return R, which does not vary across z, while the per-unit cost of labor is $\frac{w_z}{\psi_z}$, where w_z is as defined in Eq. (4) and $1/\psi_z$ is the relative productivity of labor compared to capital.

As firms minimize production costs and $k_z(0) = 0 < l_z(0)$, the above implies the existence of a threshold q^* such that tasks with complexity below q^* will be automated, while those above will use labor. The threshold q^* will be determined endogenously, as the complexity level for which it costs the same to use either capital or labor:

$$q^* = \left(\frac{w_z/\psi_z}{R}\right)^{\frac{1}{\lambda}} \tag{9}$$

The resulting cost of producing each task will therefore be:

$$p_z(q) = \begin{cases} RZ_z q^\lambda & \text{if } q \in [0, q^*] \\ \frac{w_z}{\psi_z} Z_z & \text{if } q \in (q^*, \infty). \end{cases} \tag{10}$$

Figure 9 provides a schematic diagram of the task allocation problem. Note that, as w_z varies with z, the threshold $q^*(z)$ will also vary across intermediate goods.

From the thresholds $q^*(z)$ found above, we can then integrate over the task space to find the factor shares. To do so, we assume that tasks in each sector are Pareto distributed with the following probability density function:

$$g_z(q) = \xi q^{-\xi-1}$$

where $0 < \xi < \lambda$.

We can simplify by assuming $Z_z = \frac{\lambda-\xi}{\lambda}$ and eliminating q^* in (6), which gives us:

$$Y_z = \left(\frac{k_z}{\xi/\lambda}\right)^{\frac{\xi}{\lambda}} \left(\frac{\psi_z l_z}{1-(\xi/\lambda)}\right)^{1-\frac{\xi}{\lambda}} \tag{11}$$

From this we can see that the sectoral capital share $\alpha_z = \frac{\xi}{\lambda}$. See Appendix A for a full derivation.

Figure 9: Endogenous Capital Share

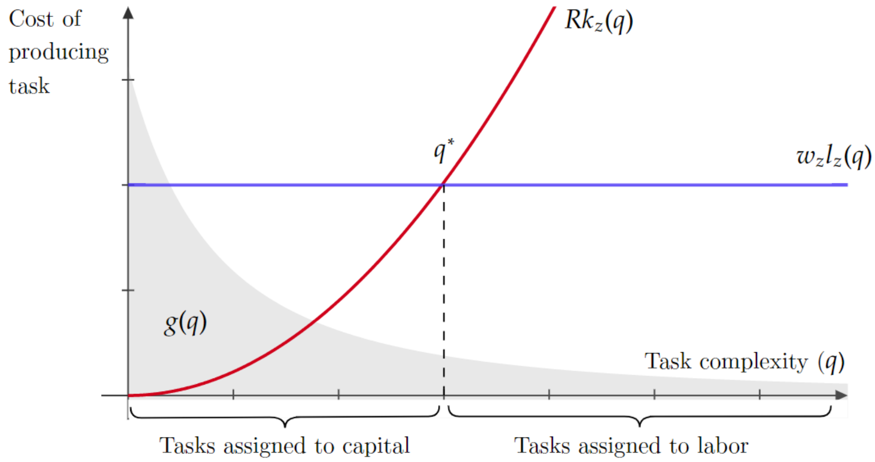

Notes: $Rk_z(q)$ represents the cost of producing a task of complexity q with capital, and $w_z l_z(q)$ represents the cost of producing it with labor. The distribution of task complexity is given by $g(q)$.

5.2 Endogenizing Technology Adoption

Imagine now that a technology exists such that for each task complexity q a firm can choose to pay $b \cdot n$ to adopt a technology that reduces the capital requirement for that task by $\kappa(n)$, where n represents the intensive margin of adoption. Examples of the intensive margin can be the number or power of machines in an industrial plant in the case of automation, or more powerful models in the case of AI. In this case, the capital requirement for a task of complexity q if the new technology is not adopted will be:

$$k_z(q) = Z_z q^\lambda$$

of if adopted, the capital requirement will be:

$$k_z(q) = \kappa(n) Z_z q^\lambda + bn$$

Where $\kappa(n) = \kappa_0 \beta_z^{-1} n^{-\beta_z}$ and $0 < \beta_z < \lambda - 1$, $\kappa_0 > 0$, which yield $\kappa < 1$.

Firms therefore face adoption decisions along two margins: an intensive one, and an extensive one.

Conditional on adopting the technology at the extensive margin, the intensive margin decision is given by:

$$\min_{n \geq 0} \kappa(n) Z_z q^{\lambda} + bn$$

subject to:

$$\kappa(n) = \kappa_0 \beta_z^{-1} n^{-\beta_z}$$

From the first order condition, firms will optimally choose:

$$n^* = \left(\frac{1}{b} \kappa_0 Z_z q^{\lambda} \right)^{\frac{1}{1+\beta_z}} \tag{12}$$

Conditional on optimal intensive margin adoption, firms also have an extensive margin adoption decision to make:

$$\min(Z_z q^{\lambda}, \kappa(n^*) Z_z q^{\lambda} + bn^*)$$

Plugging in our expression for n^* and equalizing the two arguments, the threshold q_{min} below which the firm would not adopt the technology for tasks of that complexity is given by

$$Z_z q_{min}^{\lambda} = \kappa_0^{1/\beta_z} b (\beta_z^{-1} + 1)^{\frac{1+\beta_z}{\beta_z}} \tag{13}$$

Figure 10 illustrates how firms' adoption decisions will depend on the types of tasks they perform and the relative efficiency and cost of labor versus capital. Figure 10a shows that if the costs of adoption are such that it is only cost-saving for very complex tasks, firms may prefer not to adopt the technology but to instead continue to use labor for these tasks. If the extensive adoption decision threshold $q_{min} > q^*$ there will be no adoption and therefore no change in the capital share.

However, if instead $q_{min} < q^*$, firms will adopt the new technology, as shown in Figure 10b. But they will not do so for all tasks. Intuitively, tasks that are already extremely cheap and easy to perform with capital benefit little from the expensive adoption of a new technology. In the context of AI, tasks that are already simple to automate without generative AI would not warrant investment in a new technology. But for more complex tasks, those between q_{min} and q^*, the firm will adopt the new technology. This will increase q^* and therefore the capital share, where we now have:

$$q^* = \max\left(\left(\frac{w_z}{R}\right)^{\frac{1}{\lambda}}, b^{\frac{-\beta_z}{\lambda}}\left(\frac{w_z}{R}\right)^{\frac{1+\beta_z}{\lambda}}\right)$$

Our extended model would therefore predict that the greatest adoption will occur in sectors with high benefits of adoption, which is controlled by the parameter β_z, and where there are a lot of marginal tasks that are currently cheaper to perform with labor, but that would be cheaper to perform with capital conditional on adopting a new technology. In these cases we would expect tasks of moderate complexity to drive technology adoption – those that are currently difficult to perform with capital, but which are not so complex that labor will always dominate. We can therefore think of β_z as capturing the exposure of different occupations to a new technology in this endogenous adoption setup.

Figure 10: Endogenous Technology Adoption for Tasks of Different Complexities

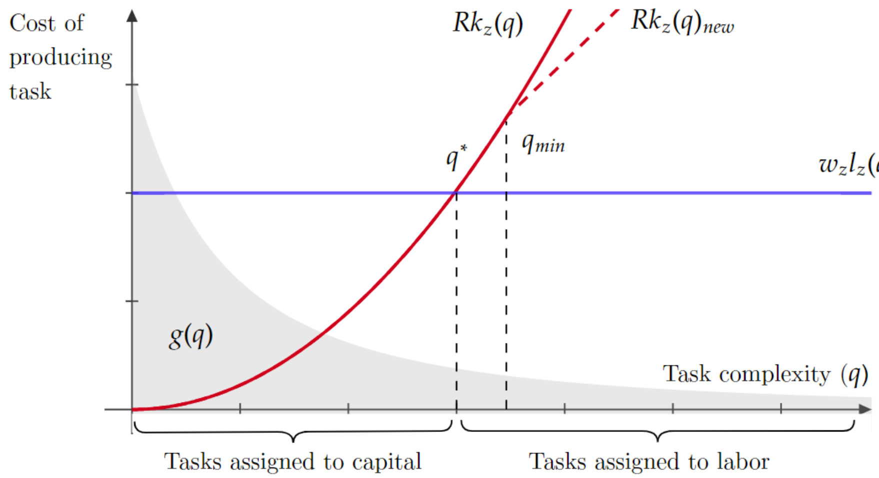

(a)No Adoption

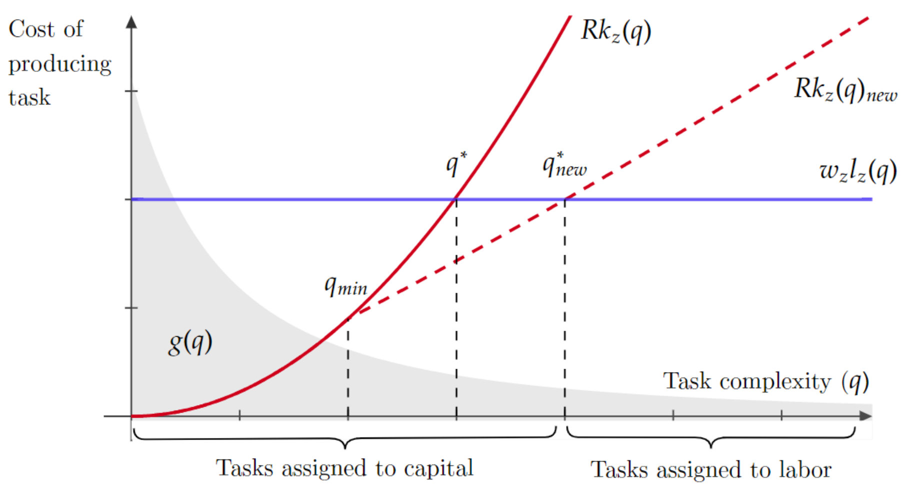

(b)Adoption for $q_{min} < q < q^*_{new}$

Notes: $Rk_z(q)$ represents the cost of producing a task of complexity q with capital, and $w_z l_z(q)$ represents the cost of producing it with labor. The distribution of task complexity is given by $g(q)$.q_{min} represents the minimum task complexity for which adopting the new technology lowers costs by more than the cost of adoption.

5.3 Calibration of Extended Model

The calibration of our extended model with endogenous capital share and technology adoption proceeds in two stages. First, we calibrate the economic environment parameters using historical data on automation from 1980 to 2014. Then, we use these calibrated parameters to model the potential impact of AI adoption from 2014 to 2048.

We begin by calibrating the task complexity distribution parameter λ and the Pareto shape parameter ξ to match the observed aggregate capital share in the UK in 1980, prior to significant automation adoption. These parameters jointly determine the initial distribution of tasks between capital and labor. Next, we model the differential impact of automation across sectors using the exposure measure ω_z, based on our Routine Task Index. The benefits of adoption for each sector, β_z, are then calculated as a function of this exposure: $\omega_z \nu$, where ν is a scalar parameter. We assume that in 1980, automation technology existed but had not yet been adopted due to high costs. We calibrate the initial adoption cost $b_{auto,1980}$ such that there is initially no adoption, but the marginal firm is indifferent between adopting and not adopting. To capture the evolution of adoption costs over time, we set the final adoption cost $b_{auto,2014}$ to be 20 percent of the initial cost, reflecting the 80 percent decline in the quality-adjusted price of industrial robots in the UK from 1980 to 2014, as reported by Graetz and Michaels (2018). Finally, we normalize the technology scalar κ_0 to 1 and calibrate the exposure scalar ν to match the observed change in the aggregate capital share in the UK between 1980 and 2014.

For the AI scenario, we maintain the same economic environment parameters (κ_0, λ, ζ, and ν) as calibrated for the automation scenario. We are implicitly assuming that the economic environment remains unchanged between the two scenarios, bar the existence of a new AI technology.

To reflect this new technology, we use a different exposure measure ω_z based on the measure of occupational exposure to AI in Felten et al. (2021). This results in a different distribution of adoption benefits β_z across sectors, where these are proportional to the exposure shown in Section 3. We calibrate the initial AI adoption cost $b_{ai,2014}$ to match the observed aggregate capital share in the UK in 2014. To model a comparable magnitude of technological change, for our scenario we assume the same percentage decline in adoption costs for AI as observed for automation. Thus, we set $b_{ai,2048}$ to be 20 percent of $b_{ai,2014}$.

This calibration allows us to simulate the potential effects of AI adoption given the same relative decline in adoption costs as observed during the automation wave of 1980-2014. All other parameters are calibrated as in the baseline model. The full set of calibrated parameters for the extended model is presented in Table 3. By using this two-stage calibration approach, we can compare the potential impacts of AI adoption to the historical effects of automation while accounting for the endogenous responses of firms in technology adoption and capital use.

Table 3: Calibration of Parameters of Extended Model

Parameter:	**Description**	**Value**	**Source**
κ_0	Technology scalar	1	Free parameter
λ	Technology efficiency	3	Drodz et al. (2022)
ζ	Task complexity dist.	1.086	Target capital share in 1980
β_z	Benefits of adoption	$\nu\omega_z$	Match exposure to automation and AI
ν	Exposure scalar	0.203	Target change in aggregate capital share
$b_{auto,1980}$	Initial adoption cost	1.3636	Marginal firm indifferent to adoption
$b_{auto,2014}$	Final adoption cost	$0.2 * b_{auto,1980}$	Graetz and Michaels (2018)
$b_{ai,2014}$	Initial adoption cost	0.5727	Marginal firm indifferent to adoption
$b_{ai,2048}$	Final adoption cost	$0.2 * b_{ai,2014}$	Match automation

We show in Appendix C that the findings of the extended model are robust to calibrating the AI scenario to have identical initial conditions to automation. For example, if we calibrate the model to match the aggregate capital share and data aggregates for 1980 rather than 2014, we obtain very similar results in terms of the implications for inequality and adoption.

5.4 Results of Extended Model

Comparing the steady state of the extended model to that of the baseline model, both for automation and AI, yields insights that both reinforce and extend our understanding of the potential impacts of technological change on inequality. While the results for automation are largely similar to those of the baseline model, the extended model produces markedly different results for AI adoption.

For the case of automation, the extended model closely mirrors the results of the baseline model. The impact on both wages (Figure 11a) and total income (Figure 11b) are almost identical, with slightly less adoption at lower incomes offset by slightly more adoption at higher incomes. This similarity is partly by design, as both models were calibrated to match the historical data from 1980-2014, constraining aggregate adoption. However, the real power of the extended model becomes apparent when we examine its predictions for AI adoption.

The extended model predicts significantly higher levels of AI adoption compared to both the baseline AI model and automation in the extended model. This increased adoption stems from a crucial mechanism captured by the extended model: the greater attractiveness of automating tasks performed by higher-paid workers. Since AI is particularly adept at tasks typically carried out by higher-wage (and thus more expensive) workers, it becomes a much more appealing technology for firms to adopt.

This enhanced adoption leads to substantial differences in outcomes, as evidenced in Table 4. The extended model predicts a much larger change in the capital share for AI adoption (10.2 p.p.) compared to both the baseline AI scenario (5.5 p.p.) and the automation scenarios. This increased capital utilization translates into larger gains in productivity and output. The extended AI model projects TFP growth of 1.7 percent and output growth of 20.7 percent, significantly higher than the baseline AI scenario (1.4 percent and 10.6 percent respectively) and the automation scenarios. However, because a lot of the higher output accrues to capital, aggregate compensation to labor changes very little, with average wages almost unchanged in all scenarios.

Figure 11: Predicted Impact of Automation on Wages and Total Income, 1980-2014

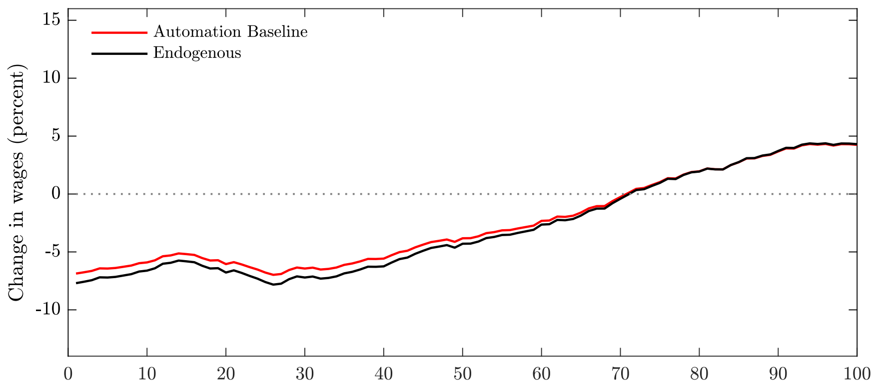

(a) Change in Wages by Income Percentile

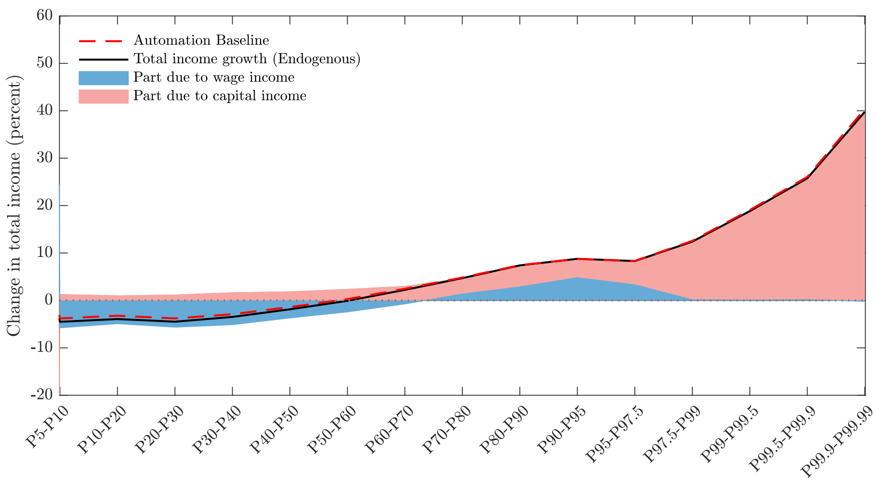

(b) Change in Total Income by Income Percentile

Notes: Figure (a) shows the percent change in wages between 1980 and 2014 by total income percentile for the baseline automation scenario in red, relative to the results of the extended model in black. Figure (b) shows the percent change in total income between 1980 and 2014 by total income percentile in the baseline automation scenario in red, relative to the results of the extended model in black, decomposed into the change due to labor income in blue, and the change due to capital income in red. Note that the scale for Figure (b) is nonlinear, with the top-tail of the income distribution expanded.

The amplified adoption in the extended model also intensifies the impact of AI on inequality. As illustrated in Figure 12a, low-income workers are expected to see even larger wage increases due to the more substantial productivity and output growth. Conversely, higher-income workers are likely to experience more significant wage decreases due to greater task displacement. This effect is reflected in the wage Gini coefficient, which is predicted to fall by 3.91 p.p. in the extended AI scenario, compared to a 1.73 percentage point decrease in the baseline AI model (Table 4).

However, as in the baseline model, the story of inequality is more nuanced when we consider total income. The greater AI adoption in the extended model leads to an even larger increase in the return on capital. As shown in Figure 12b, this results in substantially higher capital income for the highest-income households. Consequently, the wealth Gini is projected to increase by 13.67 percentage points in the extended AI scenario, nearly double the 7.16 percentage point increase predicted by the baseline AI model (Table 4).

In summary, by capturing the endogenous nature of technology adoption, our extended model reveals a more pronounced impact of AI on inequality. While it predicts a larger reduction in wage inequality due to increased productivity benefits for low-income workers, it simultaneously forecasts a more dramatic increase in wealth inequality driven by higher returns to capital. It also predicts potentially greater labor market implications than seen for previous episodes of automation due to the potential cost-savings of displacing expensive high-income workers. These results underscore the importance of considering the endogenous adoption decision both when considering the possible impacts of new technologies, and possible policy responses to them. They also highlight a dilemma for policy makers – to a far greater extent than previously seen for automation, the downsides of widespread AI adoption in the form of higher wealth inequality are likely to be accompanied by higher wages for low income workers and higher productivity.

Figure 12: Predicted Impact of AI on Wages and Total Income, 2014-2048

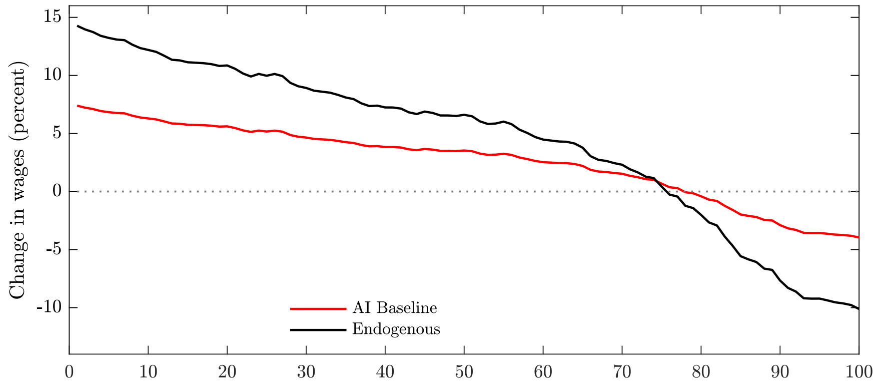

(a) Change in Wages by Income Percentile

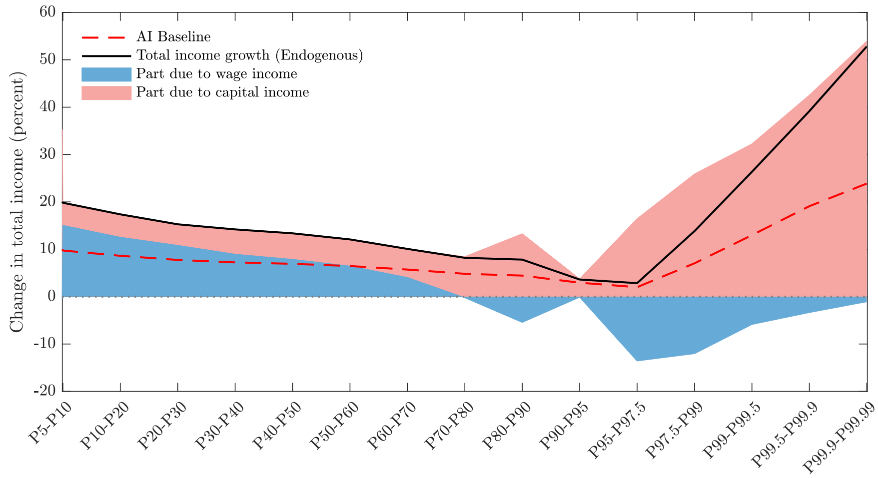

(b) Change in Total Income by Income Percentile

Notes: Figure (a) shows the percent change in wages between 2014 and 2048 by total income percentile for the baseline AI scenario in red, relative to the results of the extended model in black. Figure (b) shows the percent change in total income between 1980 and 2014 by total income percentile in the baseline AI scenario in red, relative to the results of the extended model in black, decomposed into the change due to labor income in blue, and the change due to capital income in red. Note that the scale for Figure (b) is nonlinear, with the top-tail of the income distribution expanded.

Table 4: Change in Inequality Metrics, Aggregate Productivity and Output: Extended Model

Scenario:	**Auto (baseline)**	**Auto (extended)**	**AI (baseline)**	**AI (extended)**
Wage Gini	2.05 p.p.	2.16 p.p.	-1.73 p.p.	-3.91 p.p.
Wealth Gini	6.89 p.p.	6.92 p.p.	7.18 p.p.	13.67 p.p.
Capital Share	5.5 p.p.	5.5 p.p.	5.5 p.p.	10.2 p.p.
Output	9.6%	9.6%	10.6%	20.7%
Mean Wages	0.2%	0.1%	0.2%	-0.5%

6 Efficiency Inequality Trade-Off

Now that we have endogenized the capital share and the adoption of new technologies, we can meaningfully examine the trade-off policy makers face between reducing inequality and maximizing efficiency in this setting.

As we have shown in the previous section, unfettered adoption of these new technologies is likely to lead to significant wealth inequality. However, a trade-off arises because policies aiming to reduce inequality by limiting adoption (or taxing the gains of adoption) will also reduce output by inefficiently lowering the adoption of more productive technologies and inducing misallocation between labor and capital. An additional interesting nuance in the case of AI is that high adoption could lead to lower wage inequality, further confounding the trade-off.

6.1 Implementing a Capital Tax

A natural policy response to the rising wealth inequality we document above would be to implement taxes that directly tax the capital income driving this wealth inequality. This could take the form of a capital tax, which we consider here. Such taxes may reduce inequality both directly, by lowering capital returns, and indirectly, by discouraging technology adoption. However, they also induce misallocation by distorting firms' input choices between capital and labor.

Other tax instruments could also be used to potentially address rising inequality from AI adoption. Progressive labor income taxes, for instance, could also be used to redistribute income without distorting the return on capital, potentially offering efficiency advantages over capital taxation. However, they wouldn't directly address the growing disparity in capital returns that drives much of the wealth inequality in our model, and could also introduce distortions in the labor market. A wealth tax represents another alternative that would target accumulated wealth rather than capital income flows. While our model suggests that any policy that raises the effective cost of capital would be qualitatively similar to the capital taxes we study in reducing AI adoption and associated inequality, a more detailed model with heterogeneous assets, returns, and entrepreneurial abilities would be needed to fully capture the differences between these policy instruments.[9] These considerations highlight the complex tradeoffs policymakers face when designing redistribution policies in response to technological change, though a full optimal tax analysis is beyond the scope of this paper.

In our model, a capital tax τ_K raises the effective cost of capital to $(1+\tau_K)R$. This affects firms' decisions through multiple channels, as shown in Table 5. First, by raising the cost of capital relative to labor, it reduces the threshold q^* below which firms use capital instead of labor to perform tasks. This creates misallocation even for firms that don't adopt the new technology. Second, it reduces the intensive margin of technology adoption n^* by making capital investment less attractive. Finally, although it does not directly affect the adoption threshold q_{min}, by lowering q^* it will also indirectly lower adoption at the extensive margin. Together, these effects not only lower the adoption of productivity-enhancing technologies, but also induce misallocation among non-adopting firms by distorting their capital/labor input mix (Hsieh and Klenow, 2009; Restuccia and Rogerson, 2008).

[9] While our model does not fully capture the nuanced differences between capital income taxes and wealth taxes, we can think of them as qualitatively similar in terms of their impact on equilibrium capital use and technology adoption, as both would increase the effective cost of capital and potentially discourage wealth accumulation. However, it's important to note that these taxes operate differently in practice. As discussed in Hebous et al. (2024), wealth taxes are equivalent to taxing a fixed rate of return rather than actual capital income flows. This means wealth taxes treat investments with different yields differently - high-yielding investments face relatively lighter taxation (as a percentage of income), while low-return or loss-making investments bear heavier tax burdens. If modeled explicitly, these distinctions might lead to different impacts on AI adoption patterns across sectors with varying risk-return profiles.

Table 5: Impact of Tax on Capital Use and Technology Adoption

Decision	**Tax Impact**
Capital Share	$q^* = \max\left(\left(\frac{w_z}{(1+\tau_K)R}\right)^{\frac{1}{\lambda}}, b^{\frac{-\beta_z}{\lambda}}\left(\frac{w_z}{(1+\tau_K)R}\right)^{\frac{1+\beta_z}{\lambda}}\right)$
Extensive Margin	$Z_z q_{min}^{\lambda} = \kappa_0^{1/\beta_z} b(\beta_z^{-1}+1)^{\frac{1+\beta_z}{\beta_z}}$
Intensive Margin	$n^* = \left(\frac{1}{b(1+\tau_K)R}\kappa_0 Z_z q^{\lambda}\right)^{\frac{1}{1+\beta_z}}$

6.2 Policy Simulations

To quantify the impact of the capital tax considered above on output and inequality, we introduce a government sector that collects tax revenue and redistributes it to households in the form of a Universal Basic Income (UBI).

The government's budget constraint is given by:

$$T = \tau_K RK$$

where T is the total tax revenue and K is the aggregate capital stock. Assuming a balanced budget, total government transfers to households in the form of UBI are also equal to T.

We also modify the household's budget constraint to include the UBI (represented as a lump-sum transfer each period) and the tax on capital income:

$$da_{z,t} + db_{z,t} = (r_K a_{z,t} + r_B b_{z,t} + w_z - c_{z,t} + \frac{T}{N})dt + a_{z,t} v dW_t$$

where N is the total number of households and T/N is the per-period lump sum transfer that accrues to each household.

We consider the impact of the tax in the case of automation and AI in turn. We impose a 15 percent tax rate, which would raise tax revenue equal to 4.8 percent of pre-tax GDP under automation, and equal to 5.6 percent of pre-tax GDP under AI.[10] For both automation and AI, we examine what the expected impact of each tax would be on inequality, the capital share, output, and wages, relative to the final stead states presented above absent any policy interventions.

Table 6: Tax Impacts on Inequality, Aggregate Productivity and Output

Scenario:	τ_K: **Auto**	τ_K: **AI**
Wage Gini	-0.34 p.p.	0.70 p.p.
Wage + UBI Gini	-3.83 p.p.	-3.44 p.p.
Wealth Gini	-2.12 p.p.	-3.74 p.p.
Capital Share	-0.80 p.p.	-1.60 p.p.
Output	-15.5%	-26.9%
Mean Wages	-8.7%	-11.8%

Table 6 summarizes the results of these policy simulations, showing the impact of the proposed tax and UBI on a range of aggregate outcomes. For automation, the capital tax reduces the wage Gini by 0.34 percentage points pre-transfer, and by 3.83 percentage points after accounting for the impact of UBI. As a result, the wealth Gini falls by 2.12 percentage points. However, these reductions in inequality come at a substantial cost: output falls by 15.5 percent and mean wages decline by 8.7 percent. Note that the decline in output is larger than the increase in output due to automation. That reflects the dual impact of the capital tax on output - it not only suppresses output through reduced technology adoption, it also lowers the output of non-adopting firms by inducing misallocation by distoring their input mix away from capital.

Figure 13 illustrates how these aggregate effects are distributed across workers. For automation (Figure 13a), the capital tax reduces pre-transfer wages across the distribution, to a largely similar extent. Once accounting for UBI, the lowest-income households are ultimately better off. But the average post-transfer wage remains lower than the average pre-tax pre-transfer wage.

[10]At this tax rate, we are still on the increasing side of the Laffer curve in both scenarios.

Figure 13: Changes in Pre and Post-Transfer Wages Following Capital Tax

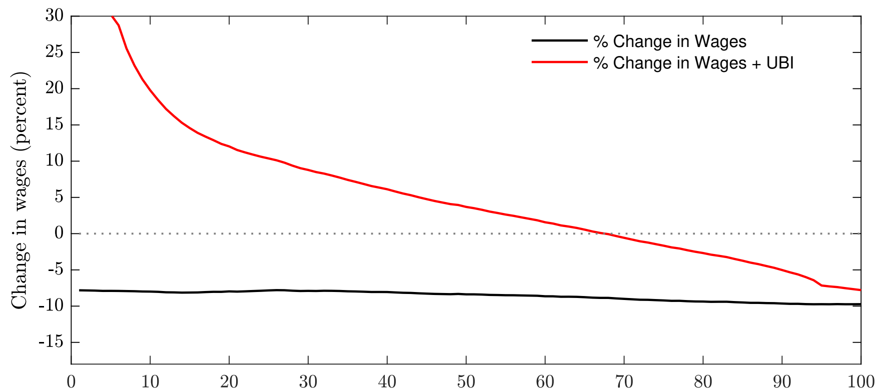

(a) Automation

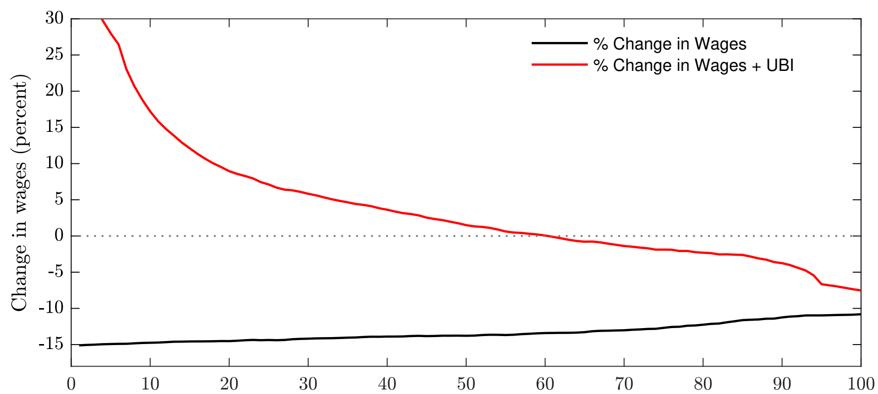

(b) AI

Notes: Both figures show the percent change in wages in the presence of the capital tax relative to no capital tax. The black line shows the change in wages not accounting for the uniform transfers, and the red line shows the change in post-transfer wages. Figure (a) shows this exercise for our automation scenario, and Figure (b) shows the results for the AI scenario.

The trade-off becomes even starker in the case of AI. While the wealth Gini falls by 3.74 percentage points, wage inequality actually increases pre-transfer (rising 0.70 percentage points) as the tax particularly dampens the adoption of AI in high-wage occupations where it would have had the largest productivity benefits. Although the UBI ultimately reduces post-transfer wage inequality by 3.44 percentage points, the economic costs are severe - output falls by 26.9 percent and mean wages decline by 11.8 percent. Figure Figure 13b shows that the pre-transfer wage effects are most negative for low-income workers, as reduced AI adoption particularly limits productivity gains in their sectors. Again the lowest income workers ultimately benefit from the UBI, but even after transfers, the average worker is substantially worse off due to the large decline in output.

These results highlight the difficult trade-off policymakers face, particularly in the context of AI adoption. While capital taxes can help reduce wealth inequality and protect the lowest-income households through redistribution, they come at a significant cost to overall economic output and average living standards. Even though many households would have higher total incomes following the policy, the deadweight loss to society is large - the total income lost across all households exceeds the tax revenue raised by more than a factor of two. This cost appears to be even larger for AI than for previous waves of automation, highlighting the importance of considering the potential impact on adoption and efficiency of any policies aiming to mitigate the distributional impacts of these new technologies.

In practice, the pace and scale of AI adoption will likely also be influenced by political economy considerations. Labor market disruptions and widening wealth inequality could generate political resistance that slows adoption through regulatory restrictions or other policy interventions. Conversely, the productivity benefits might create constituencies that accelerate adoption. These political feedback loops could lead to adoption patterns that differ from our model's predictions and potentially affect both the efficiency and equity implications of AI.

7 Conclusion

AI is likely to profoundly affect many aspects of the economy, but its implications for inequality are highly debated. In this paper we argued that resolving this debate involves recognizing that AI may have different effects on inequality via multiple channels: AI may both decrease *wage* inequality via labor market disruption while increasing *wealth* inequality via higher capital income for wealthy households. Moreover, AI-induced aggregate productivity could increase the share of workers who are unconditionally better off without fundamentally altering the inequality implications of technological change.

In support of these diverse effects, we documented that high-income workers are much more likely to work in occupations exposed to new AI technologies. While roughly 60 percent of workers at the 90th income percentile are in an occupation where a large share of tasks can be performed by AI, at the 10th percentile only 15 percent of workers are in this situation. At the same time, in the upper decile workers have the lowest share of their total income from wages, the largest wealth holdings, and the largest share of their wealth in risky but high-return assets such as equity. As such, high-income workers are not only better placed to insure themselves against potential adverse labor market impacts from AI, but also stand to benefit most from a possible AI-driven rise in capital returns. Moreover, the high-income workers most exposed to these new technologies are also more likely to be complementary with them. This might suggest that rather than being simply displaced by these technologies, as often happened to low-income workers affected by automation, the workers exposed to this latest wave of innovation may face a lower risk of job loss.

In our structural analysis, we show that under the baseline assumption that occupational exposure to AI is associated directly with task displacement, the model predicts a decrease in the Gini coefficient for wage inequality of 1.73 p.p. Wealth inequality, however, is predicted to widen, with the wealth Gini rising 7.18 p.p. For comparison, the model calibrated to routine-biased automation produces a substantial increase in both wage and wealth inequality, with the Gini rising by 2.05 p.p. and 6.89 p.p., respectively. Allowing for complementarity between labor and AI mitigates the fall in wages expected for high-income workers, limiting the extent to which inequality may fall (if decreasing at all). Moreover, an increase in aggregate productivity unleashed by AI would partially compensate for the negative effect of task displacement, thus increasing the share of workers who are unconditionally better off from AI diffusion, and further raising capital returns.

Finally, we show in our extended model that for policymakers wishing to respond to the inequality impacts of such technological change, there is an inherent welfare trade-off between equity and aggregate productivity they must take into account. This model suggests that the ability of AI to automate high-wage tasks may lead to higher adoption rates than observed in previous waves of automation. This amplified adoption intensifies both the potential productivity gains and the inequality impacts of AI. For policymakers, this presents a starker trade-off between reducing wealth inequality and promoting productivity growth than seen with previous technologies. Our policy simulations suggest that while a 15 percent capital tax could substantially reduce post-transfer inequality, the associated costs to output and productivity would be nearly twice as large for AI as for previous waves of automation. These results highlight the importance of carefully considering the efficiency implications of redistributive policies in response to AI adoption, though we leave a full optimal policy analysis to future work.

References

Acemoglu, D. and D. Autor (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. Volume 4 of *Handbook of Labor Economics*, pp. 1043–1171. Elsevier.

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo (2022). Artificial Intelligence and Jobs: Evidence from Online Vacancies. *Journal of Labor Economics 40*(S1), 293–340.

Acemoglu, D. and P. Restrepo (2018). Artificial intelligence, automation, and work. In A. Agrawal, J. Gans, and A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda*, pp. 197–236. University of Chicago Press.

Acemoglu, D. and P. Restrepo (2019, May). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives 33*(2), 3–30.

Acemoglu, D. and P. Restrepo (2022). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica 90*(5), 1973–2016.

Albanesi, S., A. D. da Silva, J. F. Jimeno, A. Lamo, and A. Wabitsch (2023). New Technologies and Jobs in Europe. Technical Report DP18220, CEPR Discussion Papers.

Alekseeva, L., J. Azar, M. Gine, S. Samila, and B. Taska (2021). The demand for AI skills in the labor market. *Labour economics 71*, 102002.

Autor, D. H. and D. Dorn (2013). The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market. *American Economic Review 103*(5), 1553–1597.

Autor, D. H., L. F. Katz, and M. S. Kearney (2006, May). The Polarization of the U.S. Labor Market. *American Economic Review 96*(2), 189–194.

Autor, D. H., F. Levy, and R. J. Murnane (2003, 11). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics 118*(4), 1279–1333.

Benhabib, J., J. Perla, and C. Tonetti (2021). Reconciling Models of Diffusion and Innovation: A Theory of the Productivity Distribution and Technology Frontier. *Econometrica 89*(5), 2261–2301.

Brazil (2019). Anuário Estatístico da Previdência Social.

Briggs, J. and D. Kodnani (2023). The Potentially Large Effects of Artificial Intelligence on Economic Growth. *Global Economics Analyst*.

Brollo, F., E. Dabla-Norris, R. de Mooij, D. Garcia-Macia, T. Hanappi, L. Liu, and A. D. M. Nguyen (2024, 6). Broadening the Gains from Generative AI. *Staff Discussion Notes 2024*(002), 1.

Brynjolfsson, E., D. Li, and L. R. Raymond (2023, April). Generative AI at Work. Working Paper 31161, National Bureau of Economic Research.

Buera, F. J. and E. Oberfield (2020). The Global Diffusion of Ideas. *Econometrica 88*(1), 83–114.

Cazzaniga, M., M. F. Jaumotte, L. Li, M. G. Melina, A. J. Panton, C. Pizzinelli, E. J. Rockall, and M. M. M. Tavares (2024). Gen-AI: Artificial Intelligence and the Future of Work. *IMF Staff Discussion Note 2024/01*.

Cazzaniga, M., C. Pizzinelli, E. Rockall, and M. Mendes Tavares (2024). Exposure to Artificial Intelligence and Occupational Mobility: a Cross-Country Analysis.

Constanzi, R. N. and C. F. dos Santos (2022). Análise dos Impactos da Reforma de 2019 sobre as Idades de Aposentadoria no Regime Geral de Previdência Social. *Informe de Previdência Social 34*(8), 5–13.

Copestake, A., M. Marczinek, A. Pople, and K. Stapleton (2023). AI and Services-Led Growth: Evidence from Indian Job Adverts.

Cortes, G. M., C. J. Nekarda, N. Jaimovich, and H. E. Siu (2020, 8). The dynamics of disappearing routine jobs: A flows approach. *Labour Economics 65*.

Dabla-Norris, E., C. Pizzinelli, and J. Rappaport (2023). Job Polarization and the Declining Wages of Young Female Workers in the United Kingdom. *Oxford Bulletin of Economics and Statistics 85*(6), 1185–1210.

DataZoom (2023). Data Zoom: Simplifying Access To Brazilian Microdata.

Drozd, L., M. Taschereau-Dumouchel, and M. Mendes Tavares (2022). Understanding Growth Through Automation: The Neoclassical Perspective.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models.

Felten, E., M. Raj, and R. Seamans (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal 42*(12), 2195–2217.

Felten, E., M. Raj, and R. Seamans (2023). How will Language Modelers like ChatGPT Affect Occupations and Industries?

Goos, M. and A. Manning (2007). Lousy and Lovely Jobs: The Rising Polarization of Work in Britain. *Review of Economics and Statistics 89*(1), 118–133.

Graetz, G. and G. Michaels (2018, 12). Robots at Work. *The Review of Economics and Statistics 100*(5), 753–768.

Hebous, S., A. Klemm, G. Michielse, and C. OsorioBuitron (2024, 1). How to Tax Wealth. *IMF How to Note 2024*(001), 1.

Hsieh, C.-T. and P. J. Klenow (2009, 11). Misallocation and Manufacturing TFP in China and India. *The Quarterly Journal of Economics 124*(4), 1403–1448.

Humlum, A. (2021). Robot Adoption and Labor Market Dynamics.

Hémous, D. and M. Olsen (2022, January). The Rise of the Machines: Automation, Horizontal Innovation, and Income Inequality. *American Economic Journal: Macroeconomics 14*(1), 179–223.

Korinek, A. and J. E. Stiglitz (2017, December). Artificial Intelligence and Its Implications for Income Distribution and Unemployment. Working Paper 24174, National Bureau of Economic Research.

Korinek, A. and J. E. Stiglitz (2018). Artificial intelligence and its implications for income distribution and unemployment. In A. Agrawal, J. Gans, and A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda*, pp. 349–390. University of Chicago Press.

Lane, M., M. Williams, and S. Broecke (2023). The impact of AI on the workplace: Main findings from the OECD AI surveys of employers and workers.

Manca, F. (2023). Six questions about the demand for artificial intelligence skills in labour markets.

Milanez, A. (2023). The Impact of AI on the Workplace: Evidence from OECD Case Studies of AI Implementation.

Moll, B., L. Rachel, and P. Restrepo (2022). Uneven Growth: Automation's Impact on Income and Wealth Inequality. *Econometrica 90*(6), 2645–2683.

OECD (2021). *Pensions at a Glance 2021.* OECD Publishing, Paris.

OECD (2023). Purchasing Power Parities (PPP) (indicator).

Perla, J., C. Tonetti, and M. E. Waugh (2021, January). Equilibrium Technology Diffusion, Trade, and Growth. *American Economic Review 111*(1), 73–128.

Pizzinelli, C., A. Panton, M. Mendes Tavares, M. Cazzaniga, and L. Li (2023). Labor Market Exposure to AI: Cross-country Differences and Distributional Implications. *IMF Working Paper 2023/216.*

Restuccia, D. and R. Rogerson (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics 11*(4), 707–720.

Ribas, R. P. and S. S. D. Soares (2008). Sobre o Painel da Pesquisa Mensal de Emprego (PME) do IBGE. *IPEA, Textos para Discussão 1348.*

Webb, M. (2020). The impact of artificial intelligence on the labor market.

Zeira, J. (1998, 11). Workers, Machines, and Economic Growth. *The Quarterly Journal of Economics 113*(4), 1091–1117.

A Model Details

A.1 Endogenizing the Capital Share

This appendix provides the detailed model derivations discussed in Section 5. For the firm's problem, we adapt the task model of Moll et al. (which follow Acemoglu and Restreppo) to endogenize the capital share. To do, we assume that the economy produces a final good, which is itself a Cobb-Douglas aggregate of perfectly competitive intermediate goods. These intermediate goods are combined from a continuum of tasks, performed by capital and labour, using a Cobb-Douglas aggregator. Our key assumption is that tasks vary in complexity (with a Pareto distribution). And while the labour required to complete a given task is flat in the complexity of the task, the capital requirements for a task are upward sloping in complexity, giving an optimal level of capital use.

A.1.1 Final Good

Households consume the final good Y, which is produced from the intermediate outputs of each sector z via Cobb-Douglas aggregation:

$$Y = A \prod_{z} Y_z^{\eta_z} \tag{14}$$

where η_z gives the share of each intermediate good in the production of the final good, and $\sum_z \eta_z = 1$.

A.1.2 Intermediate Good

Each intermediate good is itself the Cobb-Douglas aggregate of a continuum of tasks, each with complexity q:

$$\ln Y_z = \int_0^{\infty} \ln \mathcal{Y}_z(q) dq \tag{15}$$

We assume that all tasks can be produced with either capital or labour, but that low-complexity tasks are cheaper to perform with capital, and high-complexity tasks are cheaper to perform with labour, due to the capital requirement for tasks being upward sloping in complexity q:

$$k_z(q) = Z_z q^{\lambda} \tag{16}$$

$$l_z(q) = Z_z \tag{17}$$

This means that below some threshold q^* tasks will be automated, and above that threshold tasks will use labour:

$$\mathcal{Y}_z(q) = \begin{cases} k_z(q) & \text{if } q \in [0, q^*] \\ \psi_z l_z(q) & \text{if } q \in (q^*, \infty). \end{cases} \tag{18}$$

The price of producing each task will therefore be:

$$p_z(q) = \begin{cases} RZ_z q^\lambda & \text{if } q \in [0, q^*] \\ \frac{w_z}{\psi_z} Z_z & \text{if } q \in (q^*, \infty). \end{cases} \tag{19}$$

The threshold q^* will be determined endogenously, as the complexity of task for which it costs the same to produce it with either capital or labour:

$$q^* = \left(\frac{w_z/\psi_z}{R} \right)^{\frac{1}{\lambda}} \tag{20}$$

A.1.3 Factor Shares and Market Clearing

From Cobb-Douglas aggregation, we know that in equilibrium, the marginal product of capital $= \alpha_z \frac{y_z}{k_z}$ and of labour $= (1 - \alpha_z) \frac{y_z}{l_z}$. Further, because we assume the intermediate goods market is perfectly competitive and the capital and labour markets must clear, we know that capital and labour will be paid its marginal product, thus:

$$\frac{R}{p_z} = \alpha_z \frac{y_z}{k_z} \tag{21}$$

$$\frac{w_z}{p_z} = (1 - \alpha_z) \frac{y_z}{l_z} \tag{22}$$

where α_z is the capital share, R is the rental rate of capital, w_z is the wage in sector z, and p_z is the price of the intermediate good in sector z.

From the threshold q^* found above, we then also can express the factor intensities as integrals over the task space. To do so, we assume that tasks in each sector are Pareto distributed with:

$$g_z(q) = \xi q^{-\xi - 1}$$

where $0 < \xi < \lambda$.

Integrating, we can then express factor intensities as:

$$\frac{k_z}{y_z} = Z\int_0^{q^*} q^\lambda dG(q) = Z\int_0^{q^*} q^{\lambda-\xi-1}dq = Zq^{*\lambda-\xi}\frac{\xi}{\lambda-\xi} \tag{23}$$
$$\frac{l_z}{y_z} = Z\int_{q^*}^{\infty} 1dG(q) = Z\int_{q^*}^{\infty} 1^{-\xi-1}dq = Zq^{*-\xi} \tag{24}$$

Letting $Z_z = \frac{\lambda-\xi}{\lambda}$, and using our expressions for factor prices above, we have:

$$\frac{k_z}{y_z} = \alpha_z\frac{p_z}{R} = q^{*\lambda-\xi}\frac{\xi}{\lambda} \tag{25}$$
$$\frac{l_z}{y_z} = (1-\alpha_z)\frac{p_z}{w_z/\psi_z} = q^{*-\xi}\frac{\lambda-\xi}{\lambda} \tag{26}$$

Eliminating q^* from these equations gives us:

$$Y_z = \left(\frac{k_z}{\xi/\lambda}\right)^{\frac{\xi}{\lambda}}\left(\frac{\psi_z l_z}{1-(\xi/\lambda)}\right)^{1-\frac{\xi}{\lambda}} \tag{27}$$

From this we can see that $\alpha_z = \frac{\xi}{\lambda}$. From this we can further derive p_z as:

$$p_z = \alpha_z^{-1}Rq^{*\lambda-\xi}\frac{\xi}{\lambda} = Rq^{*\lambda-\xi}$$
$$p_z = R^{\frac{\xi}{\lambda}}\left(\frac{w_z}{\psi_z}\right)^{1-\frac{\xi}{\lambda}}$$

Equivalently, note that we can express the labour and capital shares as a function of q^* as:

$$\alpha_z = \frac{R}{p_z}q^{*\lambda-\xi}\frac{\xi}{\lambda} \tag{28}$$
$$(1-\alpha_z) = \frac{w_z/\psi_z}{p_z}q^{*-\xi}\frac{\lambda-\xi}{\lambda} \tag{29}$$

Finally, in order to ensure the aggregate capital market clears, we need an expression for aggregate $K = \sum_z k_z$. Following Moll et al., we normalize the price of the final good to 1, such that $p_zY_z = \eta_zY$. We can therefore express the demand for k_z in terms of aggregate output as:

$$k_z = \frac{\alpha_z p_z Y_z}{R} = \frac{\alpha_z\eta_z Y}{R} \tag{30}$$

where from above, $\alpha_z = \frac{\xi}{\lambda}$. Using $\alpha = \sum_z \alpha_z \eta_z$, we therefore have:

$$K = \alpha \frac{Y}{R} \tag{31}$$

where:

$$Y = \mathcal{A} K^{\alpha} \prod_z (\psi_z l_z)^{(1-\alpha_z)\eta_z} \tag{32}$$

where: $\mathcal{A} = A\alpha^{-\alpha} \prod_z (\eta_z(1-\alpha_z))^{-\eta_z(1-\alpha_z)} \prod_z \eta_z^{\eta_z}$

A.1.4 Endogenous Technology Adoption

We now introduce the possibility for firms to adopt a technology that reduces their capital requirements. For each task complexity q, firms can pay an adoption cost bn to reduce capital needs by $\kappa(n)$, where n represents the intensive margin of adoption (e.g., the power of AI models). The capital requirement for a task of complexity q is:

$$k_z(q) = \begin{cases} Z_z q^{\lambda} & \text{without adoption} \\ \kappa(n) Z_z q^{\lambda} + bn & \text{with adoption} \end{cases} \tag{33}$$

where $\kappa(n) = \kappa_0 \beta^{-1} n^{-\beta}$ with $0 < \beta < \lambda - 1$, $\kappa_0 > 0$, $\kappa < 1$

Firms face adoption decisions on both the intensive and extensive margins. On the intensive margin, they choose optimal technology investment by solving:

$$\min_{n \geq 0} \kappa(n) Z_z q^{\lambda} + bn$$

subject to:

$$\kappa(n) = \kappa_0 \beta^{-1} n^{-\beta}$$

From the first order condition, firms will optimally choose:

$$n^* = \left(\frac{1}{b} \kappa_0 Z_z q^{\lambda} \right)^{\frac{1}{1+\beta}} \tag{34}$$

The price of producing each task will then be:

$$p_z(q) = \begin{cases} R Z_z q^{\lambda} & \text{if } q \in [0, q^*], q < q_{min} \\ R(\kappa(n^*) Z_z q^{\lambda} + bn^*) & \text{if } q \in [0, q^*], q > q_{min} \\ \frac{w_z}{\psi_z} Z_z & \text{if } q \in (q^*, \infty) \end{cases} \tag{35}$$

Where:

$$R(\kappa(n^*)Z_z q^\lambda + bn^*) = R(\kappa_0 b^\beta Z_z q^\lambda)^{\frac{1}{1+\beta}}(\beta^{-1}+1)$$

On the extensive margin, firms decide whether to adopt by comparing:

$$\min(Zzq^\lambda, \kappa(n^*)Zzq^\lambda + bn^*)$$

Plugging in our expression for n^* and equalizing the two arguments defines the threshold q_{min} below which firms will not adopt:

$$Z_z q_{min}^\lambda = \kappa_0^{1/\beta} b(\beta^{-1}+1)^{\frac{1+\beta}{\beta}} \tag{36}$$

This implies that the threshold q^* below which firms use capital becomes piece-wise:

$$q^* = \begin{cases} \left(\frac{w_z/\psi_z}{R}\right)^{\frac{1}{\lambda}} & \text{if } q^* \leq q_{min} \\ \left(\frac{Z_z^{\frac{\beta}{1+\beta}}}{\kappa_0 b^\beta(\beta^{-1}+1)} \frac{w_z/\psi_z}{R}\right)^{\frac{1+\beta}{\lambda}} & \text{if } q^* > q_{min} \end{cases} \tag{37}$$

Let:

$$\frac{Z_z^{\frac{\beta}{\lambda}}}{\kappa_0^{1/\lambda}(\beta^{-1}+1)^{\frac{1+\beta}{\lambda}}} = 1$$

Then we can express this more compactly as:

$$q^* = \min\left(\left(\frac{w_z/\psi_z}{R}\right)^{\frac{1}{\lambda}}, b^{\frac{-\beta}{\lambda}}\left(\frac{w_z/\psi_z}{R}\right)^{\frac{1+\beta}{\lambda}}\right) \tag{38}$$

The adoption decision leads to different production functions for adopting and non-adopting firms. Let me derive each case:

A.1.5 Production Functions Under Endogenous Adoption

For Firms That Do Not Adopt

For non-adopting firms, we can express factor intensities by integrating over the task space:

$$\frac{k_z}{y_z} = Z\int_0^{q^*} q^\lambda dG(q) = Z\int_0^{q^*} q^\lambda \xi q^{-\xi-1}dq = Zq^{*\lambda-\xi}\frac{\xi}{\lambda-\xi} \tag{39}$$

$$\frac{l_z}{y_z} = Z\int_{q^*}^{\infty} 1dG(q) = Z\int_{q^*}^{\infty} 1\xi q^{-\xi-1}dq = Zq^{*-\xi} \tag{40}$$

Using $Z_z = \frac{\lambda-\xi}{\lambda}$ and our earlier expressions for factor prices:

$$\frac{k_z}{y_z} = \alpha_z \frac{p_z}{R} = q^{*\lambda-\xi}\frac{\xi}{\lambda} \tag{41}$$
$$\frac{l_z}{y_z} = (1-\alpha_z)\frac{p_z}{w_z/\psi_z} = q^{*-\xi}\frac{\lambda-\xi}{\lambda} \tag{42}$$

For Firms That Adopt

For adopting firms, factor intensities require integrating over three regions:

$$\frac{k_z}{y_z} = Z\int_0^{q_{min}} q^\lambda dG(q) + Z\int_{q_{min}}^{q^*} (\kappa_0 b^\beta Z_z q^\lambda)^{\frac{1}{1+\beta}}(\beta^{-1}+1)dG(q) \tag{43}$$
$$= Z\left(q_{min}^{\lambda-\xi}\frac{\xi}{\lambda-\xi}\right) + (\kappa_0 b^\beta Z_z)^{\frac{1}{1+\beta}}(\beta^{-1}+1)\left[(q^*)^{\frac{\lambda}{1+\beta}+\xi} - (q_{min})^{\frac{\lambda}{1+\beta}+\xi}\right]\frac{1+\beta}{\lambda+\xi(1+\beta)} \tag{44}$$

The labor intensity remains:

$$\frac{l_z}{y_z} = Z\int_{q^*}^{\infty} 1dG(q) = Zq^{*-\xi} \tag{45}$$

We can eliminate p_z by combining these expressions to get:

$$\alpha_z = \left(1 + \frac{l_z/y_z}{k_z/y_z}\frac{w_z/\psi_z}{R}\right)^{-1} \tag{46}$$

Where the ratio of factor intensities is:

$$\frac{l_z/y_z}{k_z/y_z} = \left[\frac{\xi}{\lambda-\xi}q_{min}^{\lambda-\xi}q^{*\xi} + (Z^{-\beta}\kappa_0 b^\beta)^{\frac{1}{1+\beta}}\frac{\xi}{\frac{\lambda}{1+\beta}-\xi}\frac{1+\beta}{\beta}\left(q^{*\frac{\lambda}{1+\beta}} - q_{min}^{\frac{\lambda}{1+\beta}-\xi}q^{*\xi}\right)\right]^{-1} \tag{47}$$

In both cases, this yields the production function:

$$Y_z = \left(\frac{k_z}{\xi/\lambda}\right)^{\frac{\xi}{\lambda}}\left(\frac{\psi_z l_z}{1-(\xi/\lambda)}\right)^{1-\frac{\xi}{\lambda}} \tag{48}$$

A.2 Impact of Capital Taxes

We now consider how the introduction of a capital tax affects technology adoption decisions and factor shares.

A.2.1 Capital Income Tax Effects

Let τ_K be the tax rate on capital income. This raises the effective cost of capital to:

$$R_{after-tax} = (1+\tau_K)R$$

This affects firm decisions through multiple channels:

1. **Task Assignment:** The threshold q^* between capital and labor use becomes:

$$q^* = \min\left(\left(\frac{w_z}{(1+\tau_K)R}\right)^{\frac{1}{\lambda}}, b^{\frac{-\beta_z}{\lambda}}\left(\frac{w_z}{(1+\tau_K)R}\right)^{\frac{1+\beta_z}{\lambda}}\right) \tag{49}$$

2. **Technology Adoption – Intensive Margin:** The optimal adoption intensity becomes:

$$n^* = \left(\frac{1}{b(1+\tau_K)R}\kappa_0 Z_z q^\lambda\right)^{\frac{1}{1+\beta_z}} \tag{50}$$

3. **Technology Adoption – Extensive Margin:** The threshold for adoption becomes:

$$Z_z q_{min}^\lambda = \kappa_0^{1/\beta_z} b(1+\tau_K)R(\beta_z^{-1}+1)^{\frac{1+\beta_z}{\beta_z}} \tag{51}$$

These changes affect the production cost for each task:

$$p_z(q) = \begin{cases} (1+\tau_K)RZ_z q^\lambda & \text{if } q \in [0, q^*], q < q_{min} \\ (1+\tau_K)R(\kappa(n^*)Z_z q^\lambda + bn^*) & \text{if } q \in [0, q^*], q > q_{min} \\ \frac{w_z}{\psi_z} Z_z & \text{if } q \in (q^*, \infty) \end{cases} \tag{52}$$

The resulting factor intensities become:

$$\frac{k_z}{y_z} = \alpha_z \frac{p_z}{(1+\tau_K)R} = q^{*\lambda-\xi}\frac{\xi}{\lambda} \tag{53}$$

$$\frac{l_z}{y_z} = (1-\alpha_z)\frac{p_z}{w_z/\psi_z} = q^{*-\xi}\frac{\lambda-\xi}{\lambda} \tag{54}$$

A.2.2 Government Budget and Transfers

Tax revenue funds a Universal Basic Income (UBI). The government budget constraint is:

$$T = \tau_K RK \tag{55}$$

where T is total tax revenue and K is the aggregate capital stock.

The household's budget constraint now incorporates both the capital tax and UBI transfer:

$$da_{z,t} + db_{z,t} = ((1 - \tau_K)r_K a_{z,t} + r_B b_{z,t} + w_z - c_{z,t} + \frac{T}{N})dt + a_{z,t} v dW_t \tag{56}$$

where N is the total number of households and $\frac{T}{N}$ is the per-household UBI transfer. The equilibrium wage in each sector becomes:

$$w_z = (1 - \alpha_z)\frac{\eta_z}{l_z}\mathcal{A}^{\frac{1}{1-\alpha}}\left(\frac{\alpha}{(1+\tau_K)R + \delta}\right)^{\frac{\alpha}{1-\alpha}} \prod_z (l_z \psi_z)^{\frac{\eta_z(1-\alpha_z)}{1-\alpha}} \tag{57}$$

The capital market clearing condition becomes:

$$K = \alpha \frac{Y}{(1+\tau_K)R} \tag{58}$$

This gives aggregate output under the capital tax as:

$$Y = \mathcal{A}K^{\alpha} \prod_z (\psi_z l_z)^{(1-\alpha_z)\eta_z} \tag{59}$$

where the productivity term $\mathcal{A}$ now reflects the tax distortion:

$$\mathcal{A} = A\alpha^{-\alpha}(1+\tau_K)^{-\alpha} \prod_z (\eta_z(1-\alpha_z))^{-\eta_z(1-\alpha_z)} \prod_z \eta_z^{\eta_z} \tag{60}$$

B Additional Results

B.1 Empirical Results

Figure B.1: Household Exposure to Automation and AI by Income Percentile

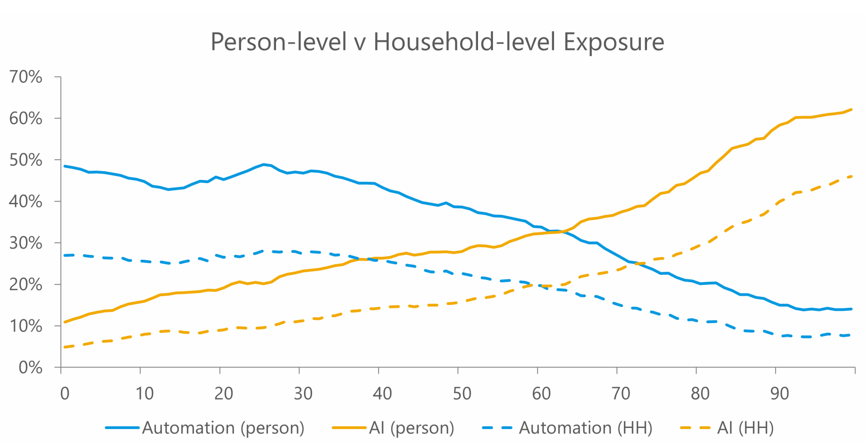

B.2 Model Fit for Automation

For the automation case, we can compare the prediction of both models to the observed change in wages in the UK between 1980 and 2014, using estimates from the ASHE (annual survey of hours and earnings). We can see that both models fit the observed change in the wage distribution relatively well. The main feature they miss is the relatively smaller decline in wages for the lowest income workers in the actual data. But this is thought to in part reflect the impact of the National Minimum Wage over this period, which would not be captured by the models.

Figure B.2: Comparison of Baseline and Extended Model Fit for Wages

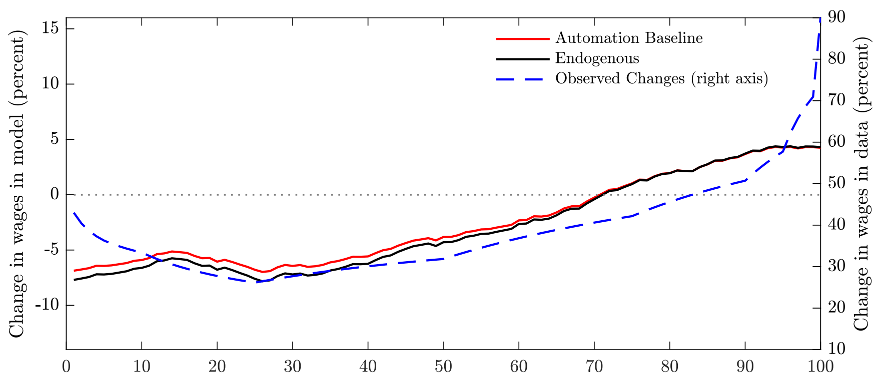

B.3 Scenario Analysis

Alternative Scenarios: Displacement and Productivity

Scenario analysis suggests the results are robust to a range of parameters. A different change in the aggregate capital share changes the 'tilt' of wage impacts, whereas a change in the aggregate productivity impact just shifts the change in wages curve higher, meaning that higher productivity is better for all, with no impact on inequality (Figure B.3).

Figure B.3: Predicted Impact of AI on Wages, 2014-2048
Alternate Paths for the Aggregate Capital Share and Productivity

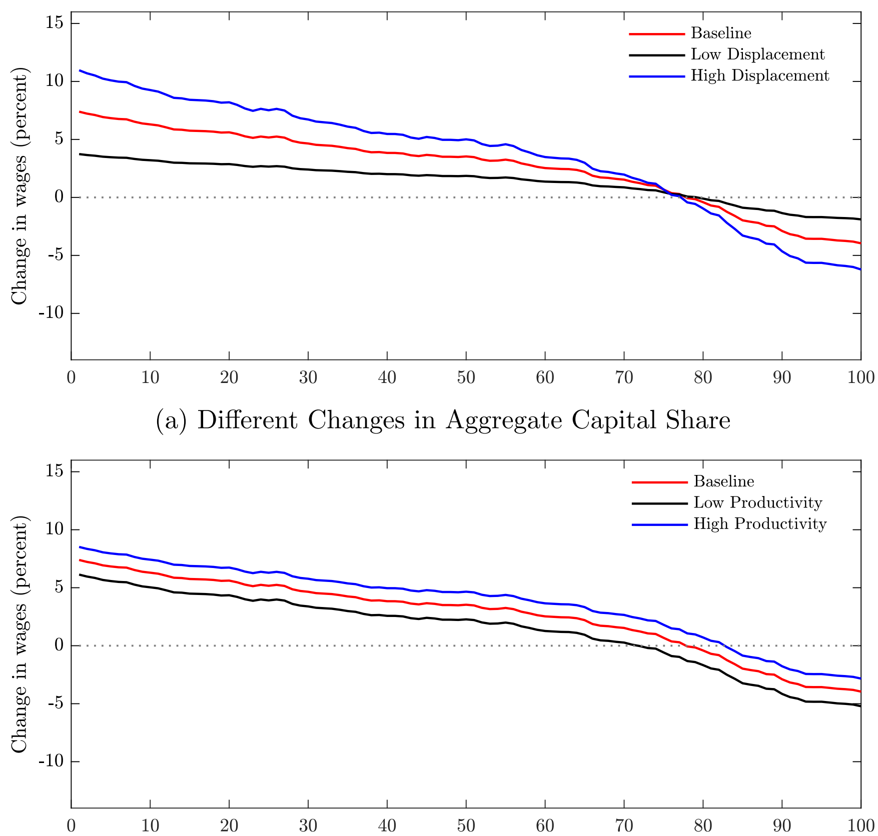

(a) Different Changes in Aggregate Capital Share

(b) Different Cost Savings From Adoption

Alternative Scenarios: Capital Supply Elasticity

In contrast, a lower capital supply elasticity leads to lower wages overall. This is because a higher cost of capital reduces output. However, the return to capital is higher in this scenario, boosting total income for the highest-income workers (Figure B.4).

Figure B.4: Predicted Impact of AI on Wages and Total Income, 2014-2048
Lower Capital Supply Elasticity

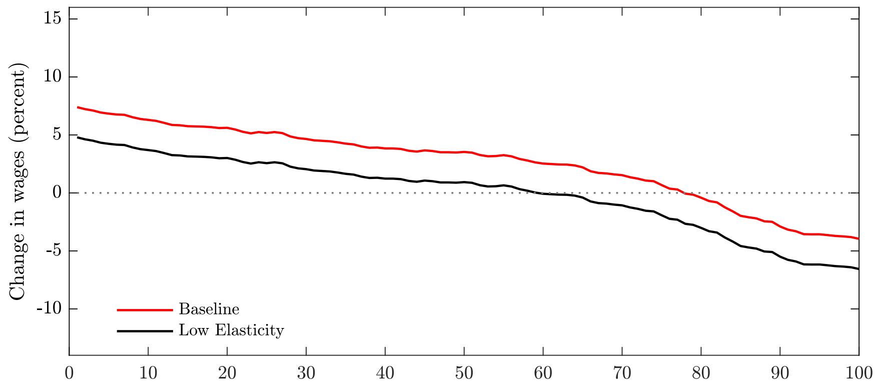

(a) Different Changes in Aggregate Capital Share

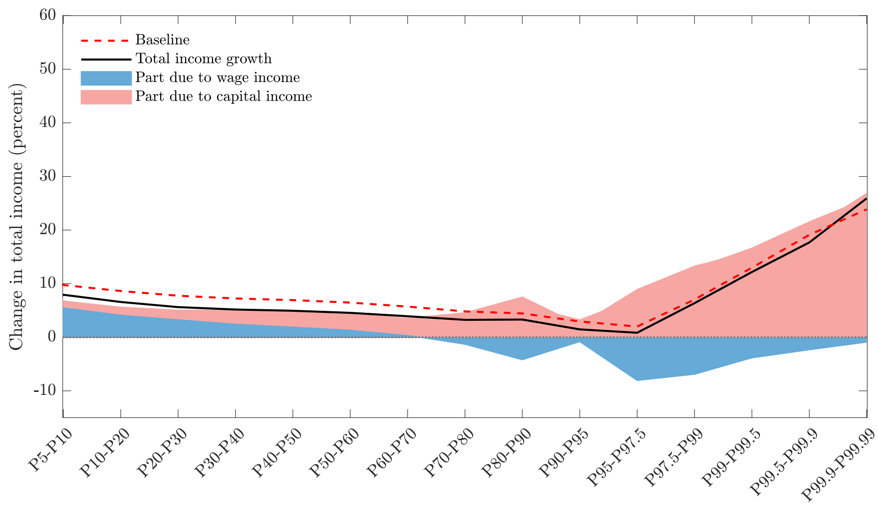

(b) Different Cost Savings From Adoption

C Calibration Robustness

Below we calibrate the extended model to match the aggregate capital share and data aggregates for 1980 rather than 2014. Doing so produces very similar results in terms of the implications for inequality and adoption, as shown in Figure C.1.

Figure C.1: Change in Wages by Income Percentile

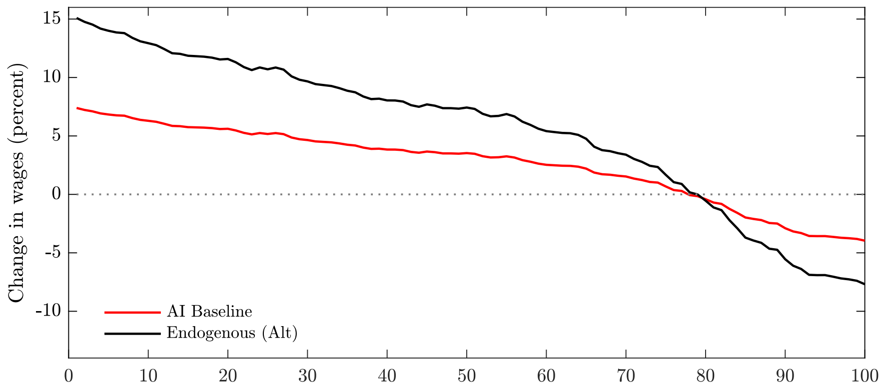

Notes: Figure C.1 shows the percent change in wages by total income percentile for the baseline AI scenario in red, relative to the results of the extended model in black, calibrated to 1980 initial conditions.