

Inefficient Automation*

Martin Beraja

MIT and NBER

Nathan Zorzi

Dartmouth

July 9, 2023

Abstract

How should the government respond to automation? We study this question in a heterogeneous agent model that takes worker displacement seriously. We recognize that displaced workers face two frictions in practice: reallocation is slow and borrowing is limited. We analyze a second best problem where the government can tax automation but lacks redistributive tools to fully alleviate borrowing frictions. The equilibrium is (constrained) inefficient and automation is excessive. Firms do not internalize that automation depresses the income of automated workers early on during the transition, precisely when they become borrowing constrained. The government finds it optimal to slow down automation on efficiency grounds, even when it does not value equity. Quantitatively, the optimal speed of automation is considerably lower than at the *laissez-faire*. The optimal policy improves efficiency and delivers meaningful welfare gains.

*Martin Beraja: maberaja@mit.edu. Nathan Zorzi: nathan.g.zorzi@dartmouth.edu. We thank our discussants, Pascual Restrepo and Henry Siu. We are grateful to Daron Acemoglu, George-Marios Angeletos, Adrien Bilal, Ricardo Caballero, Diego Comin, Arnaud Costinot, Mariacristina De Nardi, John Grigsby, Jonathon Hazell, Roozbeh Hosseini, Anders Humlum, Chad Jones, Narayana Kocherlakota, Pablo Kurlat, Monica Morlacco, Giuseppe Moscarini, Jeremy Pearce, Elisa Rubbo, Katja Seim, Aleh Tsyvinski, Gustavo Ventura, Jesús Fernández-Villaverde, Conor Walsh, Iván Werning and Christian Wolf for helpful discussions. All errors are our own.

1 Introduction

Automation technologies raise productivity but disrupt labor markets, displacing workers and lowering their earnings (Humlum, 2019; Acemoglu and Restrepo, 2022). The increasing adoption of automation has fueled an active debate about appropriate policy interventions (Lohr, 2022). Despite the growing public interest in this question, the literature has yet to produce optimal policy results that take into account the frictions that workers face in practice when they are displaced by automation.

The existing literature that justifies taxing automation assumes that worker reallocation is frictionless or absent altogether. First, recent work shows that a government that has a preference for redistribution should tax automation to mitigate its distributional consequences (see Guerreiro et al., 2022 and subsequent work by Costinot and Werning 2022; Korinek and Stiglitz 2020). This literature assumes that automation and labor reallocation are intrinsically efficient, and that the government is willing to sacrifice efficiency for equity. Second, a large literature finds that taxing capital in the long-run — and automation, by extension — might improve efficiency in economies with incomplete markets (Aiyagari, 1995; Conesa et al., 2009). This literature abstracts from worker displacement and labor reallocation.

In this paper, we take worker displacement seriously and study how a government should respond to automation. In particular, we recognize that workers face two important frictions when they reallocate or experience earnings losses. First, reallocation is slow: workers face barriers to mobility and may go through unemployment or retraining spells before finding a new job (Jacobson et al., 2005; Lee and Wolpin, 2006). Second, credit markets are imperfect: workers have a limited ability to borrow against future incomes (Jappelli and Pistaferri, 2017), especially when moving between jobs (Chetty, 2008).

We show that these frictions result in *inefficient* automation. A government should tax automation — even if it does not value equity — when it lacks redistributive instruments to fully alleviate borrowing frictions. The optimal policy *slows down* automation while workers reallocate but does not tax it in the long-run. Quantitatively, we find meaningful welfare gains from slowing down automation.

We incorporate reallocation and borrowing frictions in a dynamic model with endogenous automation and heterogeneous agents. Occupations use labor as an

input. Firms invest in automation to expand their productive capacity. Automation displaces labor as it lowers the wage of workers employed in automated occupations. Displaced workers face reallocation frictions: they receive random opportunities to move between occupations, experience a temporary period of unemployment or retraining when they do so (Alvarez and Shimer, 2011), and incur a productivity loss due to the specificity of their skills (Adão et al., 2020). Workers also face financial frictions: they are not insured against the risk that their occupation is automated and face borrowing constraints (Huggett, 1993; Aiyagari, 1994). This baseline model has the minimal elements needed to study our question. We enrich it for our quantitative analysis.

Displaced workers experience earnings losses when their occupation is automated, but expect their income to increase as they slowly reallocate and find a new job. This creates a motive for borrowing to smooth consumption during this transition. When borrowing and reallocation frictions are sufficiently severe, automated workers are pushed against their borrowing constraints.¹ Non-automated workers remain unconstrained. Firms maximize the present value of profits discounted with the equilibrium interest rate. They do not internalize that automation depresses the income of automated workers early on during the transition, precisely when they become borrowing constrained. This creates a motive for policy intervention on efficiency grounds.

In principle, the government could implement a first best if it was able to fully alleviate borrowing constraints using redistributive transfers. This is unlikely in practice, which motivates us to study second best interventions.² In particular, we analyze the constrained Ramsey problem of a government that can tax automation and implement active labor market interventions but is unable to fully alleviate the borrowing constraints of displaced workers by redistributing income.³

¹ This is consistent with the evidence. The earnings of displaced workers fall but later partially recover (Jacobson et al., 1993); including for those exposed to technological change (Braxton and Taska, 2023). Moreover, workers who lose their job indeed attempt to borrow (Sullivan, 2008), but are often unable to fully smooth consumption (Landaís and Spinnewijn, 2021) or finance their retraining (Humlum et al., 2022) while unemployed.

² Governments often do not have access to such rich instruments, which is precisely what motivates the public finance literature (Piketty and Saez, 2013). Moreover, the taxes required to pay for the transfers could tighten constraints for other workers (Aiyagari and McGrattan, 1998) and carry large dead-weight losses (Guner et al., 2021), and the take-up of transfers could be low (Schochet et al., 2012). We allow for various forms of social insurance in our quantitative model.

³ These instruments are already used in many countries. For example, US taxes vary by type of

We have two main theoretical results. Our first result shows that the equilibrium is generically constrained inefficient, as defined by [Geanakoplos and Polemarchakis \(1985\)](#), and automation is excessive. Firms fail to internalize the effects of automation on displaced workers who become borrowing constrained. This pecuniary externality is a source of inefficiency when the firm (or non-automated workers) and automated workers disagree on how they value income over time. Taxing automation and implementing active labor market interventions makes automated workers strictly better off and leaves non-automated workers indifferent — a Pareto improvement. The policy raises the income of automated workers early on during the transition, precisely when they value it more.

Our second result characterizes optimal policy for a given set of Pareto weights. To focus on the new *efficiency* rationale that we propose, we consider weights that remove any *equity* motive. These weights ensure that the government would not distort an efficient economy to redistribute.

We show that taxing automation is optimal on efficiency grounds alone. In particular, the government should *slow down* automation while labor reallocation takes place but should not intervene in the long-run. The optimal policy not only improves efficiency but also equity when the government values it. There is no trade-off, in contrast to the literature on the taxation of automation on equity grounds. As an extension, we also consider a third best problem where the government can tax automation but cannot implement active labor market interventions. This is motivated by the fact that such interventions have mixed results ([Card et al., 2018](#)) or unintended effects ([Crépon and van den Berg, 2016](#)). The rationale for taxing automation on efficiency grounds is reinforced, as borrowing constrained workers rely excessively on mobility to self-insure.

We conclude the paper with a quantitative exploration. Our goal is to evaluate the efficiency and welfare gains from slowing down automation, while allowing for various redistributive instruments. Our theoretical analysis found that the optimal policy depends on how workers value income over time, i.e., how steep their consumption profiles are. These profiles are determined by reallocation frictions and their ability to smooth consumption. Thus, we enrich our baseline model

capital and in fact *favor* automation ([Acemoglu et al., 2020](#)). South Korea reduced tax credits on automation investments, Nevada imposed an excise tax on autonomous vehicles, and the Grand Council of Geneva in Switzerland proposed to tax automated cashiers. See [Kovacev \(2020\)](#) for a detailed review.

to ensure that it performs well along these dimensions. First, we introduce idiosyncratic mobility shocks (Artuç et al., 2010), which leads to a dynamic discrete choice for reallocation and gross flows across occupations (Moscarini and Vella, 2008). Second, we add uninsured earnings risk (Floden and Lindé, 2001), which produces a realistic distribution of savings. We also allow for unemployment benefits (Krueger et al., 2016) and non-linear income taxation (Heathcote et al., 2017) to account for existing insurance that helps workers. We calibrate the model to match several key moments of the US economy. In particular, we match the dynamics of occupation-level wages since 1980 in Cortes (2016).

We find that the constrained planner slows down the speed of automation so as to increase its half-life from 16 years at the laissez-faire to 22 years at the optimum. The optimal tax reduces investments in automation especially over the first decade of the transition. The tax starts at roughly 5%, raises to 7% over a decade and then gradually declines, reaching roughly zero in year 25. Automated workers are better off — their welfare increases by 0.80% in consumption equivalent terms — whereas non-automated workers and new generations are worse off — their welfare falls by 0.19% and 0.08%, respectively. The optimal policy offsets more than half of the gap in welfare between automated and non-automated workers at the laissez-faire. Overall, the policy raises social welfare meaningfully by 0.20%.

We then consider several robustness checks and an alternative policy. First, we target a narrower definition of liquid assets. Automated workers are more likely to become borrowing constrained. They benefit more from slowing down automation and the total welfare gains increase. Second, we target a lower occupational mobility rate to reflect its decline in recent decades. The consumption of automated workers is lower than in our benchmark as they reallocate less, but the slope of their consumption profile is not meaningfully affected. Therefore, they benefit more from the intervention but the total welfare gains are mostly unchanged. Finally, as an alternative policy, we allow the government to partially insure automated workers by providing wage supplements — similar to Trade Adjustment Assistance for Workers (TAA) in the US. In present discounted terms, the government would need to give about \$20,000 to the average automated worker to deliver the same welfare gains to them as the optimal tax on automation. The aggregate fiscal cost of this policy would be orders of magnitude larger than the amount currently budgeted for TAA. This suggests that slowing down automation delivers

welfare gains that would be costly to replicate with wage supplements alone.

Our paper relates to several strands of the literature. We contribute to the literature on the labor market impact of automation (Acemoglu and Restrepo, 2018; Martinez, 2018; Humlum, 2019; Moll et al., 2021; Hémous and Olsen, 2022) by studying optimal policy in an economy with frictions and quantifying the gains from slowing down automation. Moreover, we show that taxing automation improves *both* efficiency and equity, while there is a trade-off in the efficient economies studied in the literature (Guerreiro et al., 2022; Costinot and Werning, 2022; Thuenemel, 2018; Korinek and Stiglitz, 2020).

The rationale that we propose for taxing automation also complements a large literature on capital taxation due to equity considerations (Judd, 1985; Chamley, 1986), dynamic inefficiency (Diamond, 1965; Aguiar et al., 2021), or pecuniary externalities when markets are incomplete (Conesa et al., 2009; Dávila et al., 2012; Dávila and Korinek, 2018). Optimal policies in our model also address pecuniary externalities. However, the inefficiency that we document relies neither on the presence of uninsured risk nor on endogenous borrowing constraints, contrary to the incomplete markets literature. In addition, that literature has almost exclusively studied static (or two-period) models or long-run stationary equilibria. In contrast, the rationale for intervention that we propose applies during the transition to the long run, and the *timing* of externalities is central to optimal policy.

The mechanism that we present could apply to any changes in labor demand that displace labor, including creative destruction (Caballero and Hammour, 1996) and offshoring (Hummels et al., 2018). We show that slowing down the adoption of automation technologies can improve efficiency when displaced workers are borrowing constrained. As such, our paper complements a literature studying the optimal speed of structural reforms and trade liberalization (Aghion and Blanchard, 1994; Caballero and Hammour, 1996; Neary, 1982; Mussa, 1984), which focuses on different frictions like search or investment specificity.

Methodologically, our quantitative model combines two state-of-the-art frameworks: (i) dynamic discrete choice models with mobility shocks (Artuç et al., 2010) used for studying the impact of technologies and trade; and (ii) heterogeneous agent models (Huggett, 1993; Aiyagari, 1994) used for studying consumption and insurance. Our analysis also contributes to the public finance literature using models with incomplete markets (Heathcote et al., 2017; Golosov and Tsyvinski, 2006).

2 Model

Time is continuous and there is no aggregate uncertainty. Periods are indexed by $t \geq 0$. A representative firm produces a final good by combining the output of two occupations. There is a continuum of workers with unit mass. We first describe the problem of the firm which chooses automation and labor demands. We then describe the workers' problem, including the assets they trade, the frictions they face and their sources of income. Finally, we define a competitive equilibrium.

2.1 Firm

The firm produces the final good. It combines the output of two occupations which use labor as an input. The first occupation can be automated (e.g., a routine-intensive occupation) whereas the second cannot. At time $t = 0$, the firm chooses the degree of automation α in the automatable occupation.⁴ We denote automated and non-automated occupations by $h = \{A, N\}$. At time $t \geq 0$, the firm chooses labor demands $\{\mu_t^A, \mu_t^N\}$ in both occupations.

Technology. Aggregate output is produced by combining the output y_t^h of the two occupations with a neoclassical technology

$$Y_t = G(y_t^A, y_t^N). \quad (2.1)$$

The occupations' outputs are

$$y_t^h = \begin{cases} F(\mu_t^A; \alpha) & \text{if automated } (h = A) \\ F^*(\mu_t^N) = F(\mu_t^N; 0) & \text{otherwise } (h = N) \end{cases}, \quad (2.2)$$

for some production function $F(\cdot)$ with (weakly) decreasing returns to scale in labor. Automation is labor-displacing: it decreases the marginal product of workers employed in the automated occupation.⁵ Moreover, occupations are (weak) com-

⁴ For now, automation is chosen once and for all. We introduce gradual investment later on. This allows us to clarify that the optimal policy is to *slow down* automation while labor reallocates.

⁵ It should be noted that some forms of automation might complement labor within occupations too. We focus on automation technologies that displace labor, such as industrial robots, certain types of artificial intelligence, autonomous vehicles, automated cashiers, etc.

plements, so automation increases the marginal product of the non-automated occupation.⁶ We formalize these assumptions below.

Assumption 1 (Technology). *The marginal product of workers employed in the automated occupation $\partial_{\mu^A} G(F(\mu^A; \alpha), F(\mu^N; 0))$ decreases with automation α , and the cross-partial derivative $\partial_{y^A, y^N}^2 G(y^A, y^N)$ is positive so that occupations are complements.*

Automation increases output but it comes at a cost $\mathcal{C}(\alpha)$. For example, the technology requires some continued investment due to depreciation (as in our quantitative model). We define the aggregate production function net of the cost of investing in automation

$$G^*(\mu^A, \mu^N; \alpha) \equiv G(F(\mu^A; \alpha), F(\mu^N; 0)) - \mathcal{C}(\alpha). \quad (2.3)$$

We refer to $G^*(\cdot)$ as *output* in the following.

Example. We illustrate the production function (2.3) with an example based on the model of [Acemoglu and Restrepo \(2018\)](#). The occupations operate a technology where automation and labor are perfect substitutes

$$y^A = F(\mu^A; \alpha) = \alpha + \mu^A \quad \text{and} \quad y^N = F^*(\mu^N) = \mu^N.$$

The aggregate production function is

$$G^*(\mu^A, \mu^N; \alpha) = \left[(\alpha + \mu^A)^{\frac{\nu-1}{\nu}} + (\mu^N)^{\frac{\nu-1}{\nu}} \right]^{\frac{\nu}{\nu-1}} - \mathcal{C}(\alpha),$$

where $\nu < 1$ is the elasticity of substitution across occupations.

Optimization. The firm chooses the degree of automation α and labor demands $\{\mu_t^h\}$ to maximize the value of its equity

$$\max_{\alpha \geq 0} \int_0^{+\infty} Q_t \Pi_t(\alpha) dt \quad (2.4)$$

⁶ Note that while a larger α lowers the marginal product of workers employed in the automated occupation, it can raise the *aggregate* marginal product of labor (Appendix A.7). This is the case in the quantitative model.

where $\{Q_t\}$ is the equilibrium stochastic discount factor, and

$$\Pi_t(\alpha) \equiv \max_{\mu^A, \mu^N \geq 0} G^*(\mu^A, \mu^N; \alpha) - \mu^A w_t^A - \mu^N w_t^N \quad (2.5)$$

are profits given wages $\{w_t^h\}$ and the price of the final good (normalized to 1).

We impose a regularity condition that ensures that automation is positive and finite in equilibrium. This is needed for a meaningful discussion of automation.

Assumption 2 (Interior solution). *The production function $G^*(\mu^A, \mu^N; \alpha)$ is concave in α and satisfies $\partial_\alpha G^*(\mu^A, \mu^N; \alpha)|_{\alpha=0} > 0$ and $\lim_{\alpha \rightarrow +\infty} \partial_\alpha G^*(\mu^A, \mu^N; \alpha) < 0$ for any $0 < \mu^A \leq \mu^N \leq 1$ and $\mu^A + \mu^N \leq 1$.*

2.2 Workers

Workers consume and save in financial assets. They supply inelastically one unit of labor and choose to reallocate across occupations.

Preferences. Workers' preferences over consumption flows $\{c_t\}$ are represented by

$$U = \mathbb{E}_0 \left[\int_0^{+\infty} \exp(-\rho t) u(c_t) dt \right] \quad (2.6)$$

for some discount rate $\rho > 0$ and some isoelastic utility $u(c) \equiv \frac{c^{1-\sigma}}{1-\sigma}$ with $\sigma > 0$.

Reallocation frictions. We assume that the process of labor reallocation is *slow*. At time $t = 0$, workers are equally distributed across occupations, so there is a mass $\mu_0^h = 1/2$ in automated and non-automated occupations. Workers are given the opportunity to reallocate to a new occupation with intensity λ . If they do so, they enter their new occupation with probability $1 - \iota$ or a temporary state of non-employment with probability ι . Workers exit non-employment at rate $\kappa > 0$, at which point they enter their new occupation. The non-employment state can be interpreted either as unemployment due to search frictions or as temporary exit from the labor force while workers retrain.⁷ Finally, we assume that workers in-

⁷ Workers' mobility decision is purely time-dependent, which delivers tractable expressions. We allow for state-dependent mobility in our quantitative model (Section 5). We also allow for one additional reason for slow labor reallocation (new generations gradually replacing older ones).

cur a permanent productivity loss $\theta \in (0, 1]$ after they have reallocated. This loss captures the lack of transferability of skills across occupations.

To retain tractability and abstract from idiosyncratic insurance considerations at this point, we assume that workers initially employed in each occupation form a large household.⁸ This allows them to achieve full risk sharing against the risks of being allowed to reallocate (at rate λ), becoming unemployed (probability ι), and exiting unemployment (at rate κ). In what follows, we refer to each large household as automated ($h = A$) or non-automated ($h = N$) workers.

Assets. We suppose that financial markets are incomplete: workers cannot trade contingent securities against the risk that their initial occupation is automated.⁹ Workers trade bonds and the firm's equity. Bonds are in zero net supply, and workers have no bonds initially. There is a unit of equity, which is initially in the hands of a competitive mutual fund that trades the same two assets. Workers hold an equal and fixed share in this fund, which rebates profits lump sum to them.

These assumptions ensure that (displaced) workers cannot self-insure by selling equity or bonds (since they do not hold any initially) or their share in the mutual fund. In practice, almost all the firm equity in the US is held by the wealthiest 10% of households (Survey of Consumer Finances, 2022) — not the typical displaced worker.¹⁰ We show in Appendix A.8 that our main results (Propositions 1 and 2) carry through when displaced workers do not hold any shares in the mutual fund and thus claim no profits.

Budget constraint. A worker's flow budget constraint is

$$da_t^h = \left(\hat{\mathcal{Y}}_t^h + \Pi_t + r_t a_t^h - c_t^h \right) dt \quad (2.7)$$

⁸ This assumption prevents an artificial dispersion in the distribution of assets and implies that a worker's reallocation history is irrelevant. We relax this assumption in our quantitative model.

⁹ We rule out complete markets for two reasons: financial markets participation is limited in practice (Mankiw and Zeldes, 1991); and workers' equity holdings are typically not hedged against their employment risk (Poterba, 2003). The absence of contingent securities is precisely what motivates the literature on the regulation of automation. The equilibrium would be efficient if workers could trade contingent securities before occupations become automated.

¹⁰ Accordingly, an alternative approach would have been to introduce a third agent (i.e., a Ricardian investor) who trades and holds equity but does not supply labor. We did not pursue this route to avoid introducing an additional agent.

with $a_0^h = 0$, where a_t^h is bond holdings, $\hat{y}_t^h$ is labor income, Π_t is the profits rebated by the mutual fund, and $r_t \geq 0$ is the return on savings. To save on notation, the budget constraint (2.7) implicitly assumes that workers only save in bonds. This is without loss of generality, as workers will be indifferent between saving in bonds or equity in equilibrium.¹¹ Labor income $\hat{y}_t^h$ is

$$\hat{y}_t^h = \begin{cases} w_t^A (1 - u_t - \tilde{\mu}_t) + (1 - \theta) w_t^N \tilde{\mu}_t & \text{if } h = A \\ w_t^N & \text{if } h = N, \end{cases} \quad (2.8)$$

where u_t and $\tilde{\mu}_t$ are the shares of automated workers who are unemployed or have become employed in the non-automated occupation, respectively. Expression (2.8) already uses the fact that, in equilibrium, non-automated workers do not reallocate. The expression also assumes that unemployed workers earn no income.¹²

Borrowing friction. Workers are subject to a borrowing constraint

$$a_t^h \geq \underline{a} \quad (2.9)$$

where the borrowing limit is $\underline{a} \leq 0$.

Optimization. The households maximize utility (2.6) by choosing consumption c_t^h , bonds a_t^h , and reallocation m_t^h , subject to the following constraints. First, they must satisfy the budget constraint (2.7) and borrowing constraint (2.9). Second, their labor income is given by (2.8). Third, workers' labor supply across occupations is consistent with their reallocation choice (or intensity) $m_t^h \in [0, 1]$, given reallocation frictions. Since only automated workers find it optimal to reallocate, in the following we use $m_t \equiv m_t^A$ and implicitly set $m_t^N = 0$. The laws of motion for the share of automated workers who are unemployed (u_t) or employed in the

¹¹ In equilibrium, there will be no arbitrage between bonds and equity — i.e., condition (2.14) holds. The reason is that (i) both assets are traded, and (ii) the borrowing constraint (2.9) applies to the sum of bond and equity holdings as in [Werning \(2015\)](#).

¹² Our quantitative model has unemployment benefits and gross flows across occupations.

non-automated occupation ($\tilde{\mu}_t$) are

$$du_t = [\lambda \iota (1 - u_t - \tilde{\mu}_t) m_t - \kappa u_t] dt \quad (2.10)$$

$$d\tilde{\mu}_t = [\lambda (1 - \iota) (1 - u_t - \tilde{\mu}_t) m_t + \kappa u_t] dt, \quad (2.11)$$

with $u_0 = \tilde{\mu}_0 = 0$. Next, we impose a regularity condition on reallocation frictions which ensures that reallocation takes place in equilibrium and output does not decrease over time.

Assumption 3 (Reallocation frictions). *The productivity loss θ is sufficiently small and the duration of unemployment $1/\kappa$ is sufficiently short that $1 - (1 - \theta)(1 - 1/\kappa) < Z^*$ for some $Z^* > 0$ defined in Appendix A.5.*

2.3 Equilibrium

Market clearing in the labor market requires

$$\mu_t^A = \frac{1}{2} (1 - u_t - \tilde{\mu}_t) \quad \text{and} \quad \mu_t^N = \frac{1}{2} (1 + (1 - \theta) \tilde{\mu}_t) \quad (2.12)$$

for all $t \geq 0$. The aggregate resource constraint is

$$G^* (\mu_t^A, \mu_t^N; \alpha) = \frac{1}{2} (c_t^A + c_t^N). \quad (2.13)$$

Finally, there is no arbitrage between bonds and equity, as workers and the (competitive) mutual fund can trade both. Thus, the firm discounts future cash-flows with the equilibrium interest rate r_t . The stochastic discount factor in (2.4) is

$$Q_t = \exp \left(- \int_0^t r_s ds \right). \quad (2.14)$$

We define a competitive equilibrium below.

Definition 1 (Competitive equilibrium). A competitive equilibrium consists of a degree of automation α ; and sequences for labor demands $\{\mu_t^h\}$, consumption and savings choices $\{c_t^h, a_t^h\}$, reallocation choices $\{m_t^h\}$, interest rate, stochastic discount factor, wages, profits and incomes $\{r_t, Q_t, w_t^h, \Pi_t, \mathcal{Y}_t^h\}$ such that: (i) automation and labor demands are consistent with the firm's optimization; (ii) consump-

tion, savings, and worker reallocation are consistent with workers' optimization; and (iii) the labor market clearing condition (2.12), the resource constraint (2.13), and the no arbitrage condition (2.14) are satisfied.

3 Equilibrium Characterization

We now characterize the laissez-faire equilibrium allocations. We begin with the allocations of labor, and consumption and savings *after* automation has occurred. We then turn to the equilibrium degree of automation.

3.1 Labor Reallocation and Incomes

Firm optimization implies that wages equal the marginal products of labor

$$w_t^h \equiv \partial_h G^* (\mu_t^A, \mu_t^N; \alpha) \quad (3.1)$$

for each $h = A, N$. Automation is labor-displacing and decreases the wage of automated workers so $w_t^A < w_t^N$.¹³ This induces them to reallocate to the non-automated occupation.¹⁴ As workers reallocate, the wedge between marginal products closes and output increases over time. The following proposition shows that automated workers reallocate until a stopping time T^{LF} when the marginal benefit of doing so is zero.

Lemma 1 (Equilibrium labor reallocation). *The equilibrium reallocation of labor is characterized by a stopping time T^{LF} until which automated workers reallocate to non-automated occupations. Formally, $m_t = 1$ for all $t < T^{LF}$ and $m_t = 0$ afterwards. The stopping time satisfies the smooth pasting condition*

$$\int_{T^{LF}}^{+\infty} \exp(-\rho t) \frac{u'(c_t^A)}{u'(c_0^A)} \Delta_t dt = 0 \quad (3.2)$$

¹³ We assume that wages are equalized absent automation. Otherwise, workers would have reallocated even before the firm automates at $t = 0$. Given the initial symmetric allocation of workers $\mu_0^A = \mu_0^N = \frac{1}{2}$, this requires that $\partial_A G^* (\frac{1}{2}, \frac{1}{2}, 0) = \partial_N G^* (\frac{1}{2}, \frac{1}{2}, 0)$.

¹⁴ Displacement is voluntary in competitive models such as ours. The firm offers a lower wage as it automates, and workers leave as a consequence. Introducing involuntary separations would require adding, for example, wage rigidities. This would add *another* source of inefficiency, obfuscating the mechanism that we highlight.

where

$$\Delta_t \equiv (1 - \theta) \left[\iota \left(1 - \exp \left(-\kappa \left(t - T^{\text{LF}} \right) \right) \right) + 1 - \iota \right] w_t^N - w_t^A \quad (3.3)$$

for all $t \geq T^{\text{LF}}$ denotes the output gains from labor reallocation, since in equilibrium $w_t^h = \partial_h G^* (\mu_t^A, \mu_t^N; \alpha)$.

Proof. See Appendix A.1. □

The flows Δ_t capture the benefits and costs of reallocation. When an automated worker reallocates, they forgo their wage w_t^A and earn no income if they become unemployed (probability ι) or $(1 - \theta) w_t^N$ if they enter the non-automated occupation (probability $1 - \iota$). As they exit unemployment at rate κ , they earn $(1 - \theta) w_t^N$ too. The laissez-faire stopping time T^{LF} trades off these benefits and costs.

To complete the characterization, labor allocations across occupations are

$$\mu_t^A = \frac{1}{2} \exp \left(-\lambda \min \left\{ t, T^{\text{LF}} \right\} \right) \quad (3.4)$$

$$\begin{aligned} \mu_t^N = & \frac{1}{2} + \frac{1}{2} (1 - \theta) \left(1 - \exp \left(-\lambda \min \left\{ t, T^{\text{LF}} \right\} \right) \right) \\ & - \frac{1}{2} (1 - \theta) \iota \frac{\lambda}{\lambda - \kappa} \exp(-\kappa t) \left(1 - \exp \left(-(\lambda - \kappa) \min \left\{ t, T^{\text{LF}} \right\} \right) \right), \end{aligned} \quad (3.5)$$

after solving the differential equations (2.10)–(2.11) and using labor market clearing (2.12).

Turning to earnings, automation drives a gap between the labor incomes of automated and non-automated workers $\hat{y}_t^A - \hat{y}_t^N < 0$. In theory, this gap could narrow or widen over time as labor reallocates, depending on third derivatives of the production function $G^*(\cdot)$. Reallocation directly closes this gap since fewer workers remain employed in the automated occupation. However, reallocation can indirectly affect equilibrium wages too when $G^*(\cdot)$ is concave in labor μ^A and μ^N . We next provide a sufficient condition that ensures that the effect of automation on the labor income gap $\hat{y}_t^A - \hat{y}_t^N$ weakens over time. In this case, the effect is larger when workers have not had the chance to reallocate yet, which is intuitive. The assumption uses the fact that this gap is approximately equal to $\mu \partial_\mu G^*(\mu, 1 - \mu; \alpha)$ at $\mu = \mu_t^A$ when unemployment spells are short (Assumption 3). We will only impose this assumption in Section 4.5.¹⁵

¹⁵ For instance, the example in Section 2.1 satisfies this assumption when evaluated in a symmetric allocation $y^A = y^N$.

Assumption 4 (Labor income gap). *The labor income gap $\mu \partial_\mu G^*(\mu, 1 - \mu; \alpha)$ has decreasing differences in (α, μ) .*

3.2 Consumption and Savings

We now show that the labor displacement induced by automation creates a motive for borrowing and that workers become borrowing constrained when reallocation and borrowing frictions are sufficiently severe.

Lemma 2 (Binding borrowing constraints). *Workers initially employed in the automated occupation ($h = A$) borrow in equilibrium. They become borrowing constrained if and only if reallocation frictions (λ, κ) and borrowing frictions $(\underline{a})$ are sufficiently severe. This is the case when the borrowing limit $\underline{a} \leq 0$ is sufficiently tight that $\underline{a} > a^*(\lambda, \kappa)$ for some threshold $a^*(\cdot)$ defined in Appendix A.2. This threshold satisfies $a^*(\lambda, \kappa) < 0$, i.e., borrowing constraints can bind, if and only if reallocation is slow ($1/\lambda > 0$ or $1/\kappa > 0$).*

Proof. See Appendix A.2. □


To understand this result, the left panel of Figure 3.1 depicts the paths of the labor incomes for workers initially employed in each occupation

$$\hat{\mathcal{Y}}_t^h = \underbrace{w_t^h}_{\text{Initial wage}} + \mathbf{1}_{\{h=A\}} \times 2 \times \left[\underbrace{\left(\frac{1}{2} - \mu_t^A \right) \times \left((1 - \theta) w_t^N - w_t^A \right)}_{\text{Reallocation gains}} - \underbrace{\left(1 - \mu_t^A - \mu_t^N - \left(\frac{1}{2} - \mu_t^A \right) \theta \right) \times w_t^N}_{\text{Unemployment loss}} \right]. \quad (3.6)$$

When reallocation is *slow*, automation decreases the income of workers displaced by automation, both directly by lowering the wage $w_t^A = \partial_A G^*(\cdot)$ in their initial occupation and indirectly through unemployment $(1 - \mu_t^A - \mu_t^N)$. This decrease is not fully persistent though. Their income rises over time as they become employed in the non-automated occupation at a higher wage $(1 - \theta) w_t^N$. Therefore, automated workers wish to borrow while they slowly reallocate.

Remark 1. *Workers displaced by automation expect their income to partially recover as they slowly reallocate. This creates a motive for borrowing.*

Figure 3.1: Laissez-faire: labor incomes and borrowing constraints



The right panel of the figure illustrates the second part of Lemma 2 in the space of reallocation frictions ($1/\lambda$) and borrowing frictions ($\underline{a}$) in the particular case where unemployment spells are short ($1/\kappa \rightarrow 0$). When the frictions are sufficiently mild, workers are never borrowing constrained, i.e., the white region in the figure. This region includes two limit cases in the literature. First, suppose that labor reallocation is *instantaneous* ($1/\lambda \rightarrow 0, 1/\kappa \rightarrow 0$) as in [Costinot and Werning \(2022\)](#). In this case, automated workers are still worse off (due to the reallocation cost θ), but income changes are permanent. Therefore, there is no motive for borrowing, and borrowing frictions are irrelevant. That is, slow reallocation is *necessary* for borrowing constraints to bind. Second, suppose that there are no borrowing frictions ($\underline{a} \rightarrow -\infty$) as in [Guerreiro et al. \(2022\)](#).¹⁶ In this case, automation still creates a motive for borrowing but workers are never constrained. As reallocation and borrowing frictions become more severe, borrowing constraints eventually bind $\underline{a} > a^*(\cdot)$, i.e., the colored region in the figure.¹⁷

Turning to consumption, automated workers are worse off and have a higher marginal utility, i.e., $u'(c_t^A) > u'(c_t^N)$. Automated workers have steeper consumption profiles when they are borrowing constrained, i.e., $u'(c_t^A)/u'(c_0^A) < u'(c_t^N)/u'(c_0^N) = \exp\left(-\int_0^t r_s ds\right)$.

¹⁶ In [Guerreiro et al. \(2022\)](#), reallocation takes place (entirely) through new generations replacing older ones. We introduce overlapping generations in Section 4.6.2 and in our quantitative model.

¹⁷ It should be noted that the threshold $a^*(\lambda, \kappa)$ is non-monotonic in its arguments. In particular, $\lim_{1/\lambda \rightarrow +\infty} a^*(\lambda, \kappa) = 0$ when workers cannot reallocate.

Evidence on displaced workers. The literature on the consequences of job loss has documented that the earnings of displaced workers initially fall and later partially recover (Jacobson et al., 1993), which is consistent with Remark 1. This holds in particular for workers who switch occupations due to technological change (Braxton and Taska, 2023). Moreover, workers who loose their job indeed attempt to borrow (Sullivan, 2008) but often cannot fully smooth consumption due to borrowing constraints (Landaís and Spinnewijn, 2021). Humlum et al. (2022) find that borrowing constraints affect the retraining decisions of unemployed workers. While we abstract from ex-ante heterogeneity across workers, our mechanism is more likely to be relevant when automation impacts workers with small liquidity buffers. For example, industrial robots, automated cashiers, or autonomous vehicles tend to displace low-to-middle income routine workers who are more likely to be hand-to-mouth. In contrast, artificial intelligence for natural language processing tends to affect higher income skilled workers who can borrow more easily.¹⁸

3.3 Automation

We now turn to the equilibrium automation choice.

Lemma 3 (Equilibrium automation). *The degree of automation α^{LF} is unique and interior, and satisfies*

$$\int_0^{+\infty} \exp\left(-\int_0^t r_s ds\right) \Delta_t^* dt = 0 \quad (3.7)$$

where

$$\Delta_t^* \equiv \partial_\alpha G^* \left(\mu_t^A, \mu_t^N; \alpha \right) \quad \text{for all } t \geq 0 \quad (3.8)$$

denotes the output gains from automation, and

$$Q_t = \exp\left(-\int_0^t r_s ds\right) = \exp(-\rho t) \frac{u'(c_t^N)}{u'(c_0^N)} \quad (3.9)$$

¹⁸ The mechanism might, in theory, also apply to *increases* in labor demand in an occupation or sector, as workers would borrow in anticipation of higher wages. However, this type of anticipatory effect is likely to be weak (Poterba, 1988). Indeed, we find in our quantitative model that workers borrow substantially more after a fall in their occupation's wage compared to an increase in the other occupation's (result available upon request).

is the equilibrium stochastic discount factor used by the firm. The output gains from automation Δ_t^* increase over time in equilibrium.

Proof. See Appendix A.3. □

The firm maximizes the present discounted value of output. No arbitrage between equity and bonds implies that the firm values cash-flows over time using the interest rate $\exp\left(-\int_0^t r_s ds\right)$, which equals the intertemporal marginal rate of substitution (MRS) of non-automated workers $\exp(-\rho t) u'(c_t^N) / u'(c_0^N)$ in equilibrium since they are not borrowing constrained.

The firm trades off the benefits and costs of automation over time, which are captured in the output gains Δ_t^* . These gains build up over time by Assumption 1. The reason is that workers reallocate from the automated occupation, where automation lowers their marginal product, to the non-automated occupation, where automation raises it.¹⁹

4 Excessive Automation

In this section, we show that automation is *excessive* at the laissez-faire and characterize optimal policy. We first specify the set of policy instruments available to the government (Section 4.1). We then state the constrained Ramsey problem (Section 4.2), and discuss how automation affects workers' welfare (Section 4.3). Next, we show that the equilibrium is constrained inefficient, and the government can Pareto improve upon the laissez-faire by taxing automation (Section 4.4). We then show that the government finds it optimal to tax automation purely on efficiency grounds (Section 4.5), even when it does not value equity. Finally, we present various extensions (Section 4.6). For tractability and to obtain more compact expressions, we assume in the following that workers cannot borrow $\underline{a} \rightarrow 0$.

¹⁹ We have assumed that the cost of automation $\mathcal{C}(\alpha)$ is constant over time. All of our results carry through with any time-varying cost $\mathcal{C}_t(\alpha)$ that decreases over time (i.e., the continuous time equivalent of a “sunk” cost). The reason is that the output gains Δ_t^* would increase over time *even more* in this case.

4.1 Policy Instruments

A government that has access to a sufficiently rich set of lump-sum transfers to fully undo borrowing frictions could, in theory, implement a first best. For example, the government could use targeted lump-sum transfers $\{T_t^h\}$ (indexed by worker and time) to help displaced workers. In practice, such rich interventions are unlikely. The literatures on optimal taxation (Piketty and Saez, 2013) and the regulation of automation precisely rule out such transfers, in part due to their informational requirements.²⁰ This motivates us to study second best policies interventions.

We assume that the government has access to a simple set of instruments that depend on calendar time alone: a linear tax on automation τ^a , and active labor market interventions (Card et al., 2018) that tax or subsidize labor reallocation $\{\zeta_t\}$.²¹ These instruments are already used in many economies and do not require the government to know which occupations are automated or which workers are displaced. For instance, US taxes vary by type of capital (e.g., equipment, software, structures) and industry (due to differential depreciation allowances), and seem to be favoring automation instead of taxing it (Acemoglu et al., 2020). Concrete policies discriminating against automation technologies (Kovacev, 2020) include: (i) South Korea’s reduction in the automation tax credit aimed at protecting workers in high-tech manufacturing; (ii) Nevada’s excise tax on transportation companies using autonomous vehicles that would displace human drivers; and (iii) the Swiss canton of Geneva’s proposed tax on retail stores installing automated cashiers. That said, identifying technologies that displace labor could be more challenging in other instances (e.g., artificial intelligence algorithms).

²⁰ Alternatively, the government could implement *symmetric* transfers $\{T_t\}$ to effectively borrow on behalf of the workers. However, the associated debt needs to be repaid later by taxing them. This future tax burden could tighten borrowing constraints (Aiyagari and McGrattan, 1998) and carry large distortions (Guner et al., 2021), limiting or reversing the benefits of the transfers. The transfers need to be generous enough to ensure that *no* worker is constrained — a scenario that the literature on heterogeneous agents has not seriously considered. The size of transfers is further limited by the fact that future higher taxes could push the poorest workers into default.

²¹ To abstract from income effects, we assume that the large families reimburse lump sum any reallocation taxes or subsidies they receive. These can take the form of credits for retraining programs or unemployment insurance (when positive), or penalties such as imperfect vesting of retirement funds (when negative).

4.2 The Constrained Ramsey Problem

We consider the problem of a government that values automated and non-automated workers, and assigns them Pareto weights $\{\eta^A, \eta^N\}$. The government effectively controls two choices with its tax on automation and active labor market interventions: the degree of automation α ; and the reallocation of workers, as governed by the stopping time T .²² All other choices must be consistent with workers' and the firm's optimality.

Lemma 4 (Primal problem). *The government maximizes the social welfare function*

$$\mathcal{U} = \sum_h \eta^h \int_0^{+\infty} \exp(-\rho t) u(c_t^h) dt \quad (4.1)$$

by choosing $\{\alpha, T, \mu_t^A, \mu_t^N, c_t^A, c_t^N\}$, subject to the laws of motion (3.4)–(3.5) for labor $\{\mu_t^A, \mu_t^N\}$, and the consumption allocations $c_t^h = \hat{\mathcal{Y}}_t^h + \Pi_t$ for workers initially employed in occupations $h = \{A, N\}$, where labor incomes $\hat{\mathcal{Y}}_t^h$ are given by (3.6) and profits Π_t are given by (2.5).

It is worth noting that the only difference between this constrained problem and the unconstrained (first best) Ramsey problem lies in the set of implementable consumption allocations. In the constrained problem, workers must consume their income, since borrowing is not possible ($\underline{a} \rightarrow 0$). In the first best problem, any consumption allocation that satisfies the resource constraint (2.13) is feasible.

4.3 Efficiency vs. Equity Motives for Intervention

Consider the effect of a policy intervention $\{\delta\alpha, \delta T\}$ on the government's objective $\mathcal{U}$ starting from the laissez-faire. This perturbation affects workers' incomes. We denote the subsequent consumption changes by δc_t^h . The change in welfare is

²² Formally, the government would control reallocation choices $\{m_t^h\}$. To save on notation, we directly impose that the optimal reallocation policy is a stopping time T for automated workers.

$$\begin{aligned}
\delta \mathcal{U} = & \eta^N \times u' \left(c_0^N \right) \times \underbrace{\int_0^{+\infty} \exp(-\rho t) \frac{u' \left(c_t^N \right)}{u' \left(c_0^N \right)} dt}_{=\exp\left(-\int_0^t r_s ds\right)} \times \delta c_t^N dt \\
& + \eta^A \times u' \left(c_0^A \right) \times \underbrace{\int_0^{+\infty} \exp(-\rho t) \frac{u' \left(c_t^A \right)}{u' \left(c_0^A \right)} dt}_{\text{How automated workers value flows}} \times \delta c_t^A dt \quad (4.2)
\end{aligned}$$

with consumption changes

$$\delta c_t^h \equiv \delta \alpha \times \left(\hat{c}_t^{h,*} + \bar{c}^{h,*} \right) + \delta T \times \left(\hat{c}_t^h + \bar{c}^h \right), \quad (4.3)$$

where

$$\hat{c}_t^{h,*} \equiv \partial_\alpha c_t^h - \bar{c}^{h,*} \quad \text{and} \quad \bar{c}^{h,*} \equiv \int_0^{+\infty} \frac{\exp(-\rho t) u' \left(c_t^N \right)}{\int_0^{+\infty} \exp(-\rho s) u' \left(c_s^N \right) ds} \partial_\alpha c_t^h dt, \quad (4.4)$$

and $\partial_\alpha c_t^h$ denotes the marginal effect of a perturbation in automation $\delta \alpha$ on worker's h consumption at time t . The *time-varying* effects $\hat{c}_t^{h,*}$ capture how the perturbation $\delta \alpha$ affects the timing of a worker's consumption. The *permanent* effects $\bar{c}^{h,*}$ capture how this perturbation affects a worker's consumption on average. We define the permanent effects $\bar{c}^{h,*}$ so that they are purely distributional $\bar{c}^{A,*} + \bar{c}^{N,*} = 0$.²³ The terms $\hat{c}_t^h$ and $\bar{c}^h$ are defined similarly for a perturbation in reallocation δT (Appendix A.4).

Equity motive. Consider first an efficient economy without borrowing constraints. The timing of workers' incomes is irrelevant in this case. The intertemporal MRS of automated and non-automated workers coincides with the equilibrium interest rate $\exp(-\rho t) u' \left(c_t^h \right) / u' \left(c_0^h \right) = \exp \left(-\int_0^t r_s ds \right)$ for $h = A, N$. As a result, a perturbation in automation $\delta \alpha$ does not affect welfare through the time-varying effects $\hat{c}_t^{h,*}$.²⁴ However, automation has permanent and purely distributional consequences: it makes automated workers worse-off relative to non-automated workers $\bar{c}^{A,*} = -\bar{c}^{N,*} < 0$. Automated workers value these permanent changes more

²³ This follows from the fact that the firm and workers are already optimizing at the laissez-faire, i.e., equations (3.2) and (3.7) hold, and the aggregate resource constraint (2.13) holds.

²⁴ Similarly, a perturbation in reallocation δT does not affect welfare through $\hat{c}_t^h$ either.

since $u'(c_0^A) > u'(c_0^N)$. This provides a motive for taxing automation when the government values equity, e.g., when it uses utilitarian weights $\eta^h \equiv 1/2$. This equity motive has been the focus of the existing literature (Guerreiro et al., 2022; Costinot and Werning, 2022).

Efficiency motive. Suppose instead that reallocation and borrowing frictions are sufficiently severe that borrowing constraints bind (Lemma 2). The timing of workers' incomes becomes relevant in this case. Automated workers are effectively more impatient $u'(c_t^A) / u'(c_0^A) < u'(c_t^N) / u'(c_0^N)$ since they are borrowing constrained. Therefore, a perturbation in automation $\delta\alpha$ now also affects welfare through the time-varying effects $\hat{c}_t^{h,*}$. In particular, automation depresses the income (and hence consumption) of automated workers early on during the transition, precisely when they are borrowing constrained. The firm does not internalize this effect on workers' incomes when it automates. This pecuniary externality is a source of inefficiency when the firm (or non-automated workers) and automated workers disagree on how they value income over time. As we show in the next sections, this creates room for Pareto improvements and a new efficiency motive for taxing automation.

4.4 Constrained Inefficiency

We now establish that the equilibrium is *generically* constrained inefficient in the sense of Geanakoplos and Polemarchakis (1985) when displaced workers are borrowing constrained. The government can implement a Pareto improvement by varying automation ($\delta\alpha$) and reallocation (δT). This is the case in virtually *any* economy: if this happens not to be the case, then there exists an arbitrarily small perturbation of the production function $G^*(\cdot)$ that again allows for a Pareto improvement. In particular, automation is excessive and the Pareto improvement requires taxing it.

Proposition 1 (Constrained inefficiency). *Generically, there exists a variation $\{\delta\alpha, \delta T\}$ starting from the laissez-faire which makes automated workers strictly better off ($\delta U^A > 0$) and non-automated workers indifferent ($\delta U^N = 0$). The Pareto improvement requires taxing automation ($\delta\alpha < 0$).*

Proof. See Appendix A.4. □

To understand why taxing automation generates a Pareto improvement, we reproduce the key steps of the proof. Consider a reduction in automation $\delta\alpha < 0$. Absent any change in reallocation, non-automated workers would be worse off. To leave them indifferent $\delta U^N = 0$, the government can always compensate them by reducing the mass of workers who enter their occupation $\delta T < 0$, which lifts their wage. However, reducing reallocation hurts automated workers. There is a Pareto improvement when the gains from less automation outweigh the losses from less reallocation for automated workers $\delta U^A > 0$. This is the case if and only if the time-varying effects satisfy

$$\int_0^\infty \left\{ \omega_t^A - \omega_t^N \right\} \times \left\{ \delta\alpha \times \hat{c}_t^{A,*} - \delta T \times \hat{c}_t^N \right\} dt > 0, \quad (4.5)$$

where $\omega_t^h \equiv \exp(-\rho t) u'(c_t^h) / \int_0^{+\infty} \exp(-\rho s) u'(c_s^h) ds$ captures how worker h values consumption at t relative to permanent consumption, i.e., how they value the time-varying effects of a policy intervention.

Reducing automation $\delta\alpha < 0$ raises the income, and hence the consumption, of automated workers ($\delta\alpha \times \hat{c}_t^{A,*} > 0$) at times when they value it more ($\omega_t^A > \omega_t^N$). Reducing reallocation $\delta T < 0$ to compensate non-automated workers ($\delta T \times \hat{c}_t^N > 0$) lowers the income of automated workers in the future when they value it less ($\omega_t^A < \omega_t^N$). The intervention Pareto improves upon the laissez-faire when the first effect dominates. In particular, Pareto improvements are *only* possible if borrowing constraints bind (since otherwise $\omega_t^A = \omega_t^N$) and the intervention affects the *timing* of consumption.

Using the reduction in reallocation $\delta T < 0$ that ensures that $\delta U^N = 0$ after a reduction in automation $\delta\alpha < 0$, the inequality (4.5) becomes

$$\int_0^{T^{LF}} \exp(-\rho t) \left\{ \frac{u'(c_t^A)}{u'(c_{T^{LF}}^A)} - \frac{u'(c_t^N)}{u'(c_{T^{LF}}^N)} \right\} \times \left\{ \delta\alpha \times \partial_\alpha c_t^A \right\} dt > 0, \quad (4.6)$$

where $\delta\alpha \times \partial_\alpha c_t^A$ for $t < T^{LF}$ captures how automation affects the income of automated workers while they reallocate. Taxing automation $\delta\alpha < 0$ thus generates a Pareto improvement because it increases the income of automated workers when they value it the most, i.e., at times $t < T^{LF}$ when $u'(c_t^A)/u'(c_{T^{LF}}^A) > u'(c_t^N)/u'(c_{T^{LF}}^N)$. The intervention partly undoes the effect of the firm's automation choice at the laissez-faire, which depresses the labor income of automated workers

early on during the transition.²⁵

Remark 2. *Firms fail to internalize the effects of automation on displaced workers who become borrowing constrained. Taxing automation increases their income early on during the transition precisely when they value it more.*

The inefficiency that we document relies on the firm not internalizing that automation depresses the income of displaced workers at times when they value it the most. In practice, the wealthiest 10% of households hold close to 90% of firm equity in the US (Survey of Consumer Finances, 2022). The typical displaced worker does not hold or trade equity. Thus, the incentives of a firm to automate (say a car manufacturer) mostly reflect the incentives of its wealthy investors rather than those of the workers that it displaces (who can become borrowing constrained). That said, our mechanism could be muted if workers were represented in the boardroom.

4.5 Optimal Policy Interventions

We now characterize the constrained efficient degree of automation for a given set of Pareto weights. The optimal policy depends on how the government values *efficiency* and *equity*.²⁶ To see this, consider the social incentive to automate $\partial_\alpha \mathcal{U}$ starting from the laissez-faire. It can be decomposed as

$$\begin{aligned} \partial_\alpha \mathcal{U} = & \underbrace{\sum_h \eta^h \int_0^{+\infty} \exp(-\rho t) u'(c_t^h) \times \hat{c}_t^{h,*} dt}_{\text{Taxing } \alpha \text{ on efficiency grounds}} \\ & + \underbrace{\sum_h \eta^h \int_0^{+\infty} \exp(-\rho t) u'(c_t^h) dt \times \bar{c}^{h,*}}_{\text{Taxing } \alpha \text{ on equity grounds}}, \end{aligned} \quad (4.7)$$

where $\hat{c}_t^{h,*}$ and $\bar{c}^{h,*}$ are the time-varying and permanent effects of automation (Section 4.3). The efficiency component captures the effect of varying the timing of

²⁵ Taxing automation also benefits displaced workers by changing the timing of their income from profits. However, the effect on labor earnings is likely to be the most relevant in practice, as we have argued in Section 2.2. It is also the main driver of the welfare gains from taxing automation in our quantitative model. Appendix A.8 shows that Propositions 1 and 2 carry through even when displaced workers claim no profits.

²⁶ Bhandari et al. (2021) provides a decomposition of the welfare effects of policy into aggregate and redistribution components. Appendix A.9 discusses how our decomposition relates to theirs.

incomes (and so consumptions) over time. This component is zero in any efficient economy where the MRS are equalized across workers. In turn, the equity component captures how consumption is redistributed across workers on average. This component depends on the difference between workers' average marginal utilities.

Efficiency motive. To focus on the new *efficiency* rationale that we propose, we first consider a government that uses weights $\eta^{\text{effic},h} = 1 / \int_0^{+\infty} \exp(-\rho t) u'(c_t^h) dt$ evaluated at the laissez-faire. These efficiency weights ensure that the government does not intervene to improve equity, since $\bar{c}^{A,*} + \bar{c}^{N,*} = 0$.²⁷ Proposition 2 below shows that taxing automation is optimal on efficiency grounds alone.

Proposition 2 (Taxing automation on efficiency grounds). *Suppose that the government uses efficiency weights $\eta^{\text{effic},h} = 1 / \int_0^{+\infty} \exp(-\rho t) u'(c_t^h) dt$. Then, taxing automation is optimal.*

Proof. See Appendix A.5. □

The intuition is similar to Proposition 1. Automation initially depresses the income (and hence consumption) of automated workers relative to their permanent level. That is, the time-varying effect of automation is negative ($\hat{c}_t^{A,*} < 0$) early on, and becomes positive ($\hat{c}_t^{A,*} > 0$) later in the transition. This increase is monotonic over time (Assumption 4). Therefore, taxing automation is optimal ($\partial_\alpha \mathcal{U} < 0$) because it raises consumption precisely at times when the average worker in the economy (under weights $\eta^{\text{effic},h}$) values it more.

Finally, when the government's problem (4.1) is convex, Proposition 2 implies that the laissez-faire level of automation is excessive compared to its second best counterpart $\alpha^{\text{SB,effic}} < \alpha^{\text{LF}}$. The reason is that there is a unique global optimum. We will compute the optimum numerically in our quantitative model.

Equity motive. Taxing automation not only improves efficiency but also equity when the government values it. There is no trade-off, contrary to the literature on the taxation of automation on equity grounds. A utilitarian government ($\eta^{\text{util},h} = 1/2$) that values equity would tax automation even more compared to Proposition 2.

²⁷ In an efficient economy, the weights boil down to the standard inverse marginal utility weights $\eta^{\text{effic},h} = 1 / u'(c_0^h)$ and the government does not intervene at all.

4.6 Extensions

We next consider a number of extensions to our analysis.

4.6.1 No Active Labor Market Interventions

In practice, ex post policies can be difficult to implement. Active labor market interventions often produce mixed results (Card et al., 2018), or have unintended consequences for untargted workers (Crépon and van den Berg, 2016). For instance, this would be the case with *gross* flows between occupations, as in our quantitative model. For this reason, we now consider a *third best* problem where the government controls automation but is unable to control labor reallocation. This implies that a Pareto improvement (Proposition 1) is no longer possible.

In addition to the direct effects in (4.7), the government now internalizes the indirect effect of automation due to reallocation $T(\alpha)$ on workers' consumption $\partial_T c_t^h$. This indirect effect is²⁸

$$T'(\alpha) \times \frac{1}{2} \lambda \exp(-\lambda T) \times \int_{T(\alpha)}^{+\infty} \exp(-\rho t) \left\{ \eta^N u'(c_t^N) - \eta^A u'(c_t^A) \right\} \times \partial_T c_t^N dt.$$

Taxing automation decreases reallocation since $T'(\cdot) > 0$. This indirect effect can either reinforce or dampen the government's incentives to tax automation, depending on the Pareto weights. For instance, a utilitarian government would tax automation *less* compared to Proposition 2, as this induces more reallocation $\delta T > 0$ and redistributes towards automated workers. In contrast, one can show that a government using efficiency weights (which does not value such redistribution) finds it optimal to tax automation *more* when unemployment spells are not too long (as in Assumption 3).

4.6.2 Slowing Down Automation

An extensive literature argues that taxing capital might improve insurance (Conesa et al., 2009; Dávila et al., 2012) or prevent capital overaccumulation (Aiyagari, 1995) in economies with incomplete markets (or overlapping generations). These two rationales share two features: they rely on the presence on uninsured idiosyncratic

²⁸ This expression uses the fact that $\Delta_t = 1/2 (\partial_T c_t^A + \partial_T c_t^N)$ together with Lemma 1.

risk and optimal policies affect investment in the *long-run*.

The rationale that we propose is conceptually distinct. First, we find that taxing automation is optimal even absent idiosyncratic uncertainty. Second, our mechanism implies that the government should *slow down* automation only while labor reallocation takes place and displaced workers are borrowing constrained, but it has no reason to tax automation in the long-run. To clarify this last point, we extend our model along two dimensions that are relevant for studying dynamics over long horizons. Both dimensions are present in our quantitative model. First, we allow for gradual investments in automation. We assume that the law of motion of automation is $d\alpha_t = (x_t - \delta\alpha_t) dt$ for some depreciation rate δ and gross investment rate x_t , and the investment cost of automation q_t declines over time. Second, we assume that there are overlapping generations of workers who are born (and die) at rate χ and can choose any occupation at birth. We show below that the government has no motive to intervene in the long-run.

Proposition 3 (No intervention in the long-run). *In the long-run, the equilibrium converges to a first best allocation. In particular, $\alpha_t^{LF} / \alpha_t^{FB} \rightarrow 1$ as $t \rightarrow +\infty$, where α_t^{FB} is automation at the first best.*

Proof. See Appendix [A.6](#). □

The government can improve neither efficiency nor equity in the long-run. Once labor reallocation is complete, workers' incomes are constant and they have no incentives to borrow. The intertemporal MRS of all workers are identical, and the firm's automation choice is efficient. Moreover, the entry of new generations equalizes wages across occupations in the long-run. The workers' marginal utilities are equalized; there is no need for redistribution as in [Guerreiro et al. \(2022\)](#).

4.6.3 The Direction of Investments

So far, the firm could only invest in automation. Taxing it thus unequivocally reduces *total* investment. We now allow investments in a Hicks-neutral technology. We assume that aggregate output is

$$\tilde{G}(\mu^A, \mu^N; \alpha, A) = AG\left(F(\mu^A; \alpha), F(\mu^N; 0)\right) - \mathcal{C}(\alpha) - \Phi(A)$$

and the firm chooses automation α and productivity A . Hicks-neutral investments do not cause worker displacement. The adjustment is instantaneous and workers are not borrowing constrained. Therefore, the optimal policy changes the *direction* of investments: it taxes automation relative to Hicks-neutral investments.

It is also worth noting that our analysis abstracts from other reasons why the government might want to subsidize investment, e.g., firm credit constraints, externalities, etc. Therefore, our results do not necessarily imply that automation should be taxed *on net*. Rather, they suggest that automation should be subsidized *less* than other investments, e.g., through lower subsidies as in South Korea.

5 Quantitative Model

In the rest of the paper, we quantitatively evaluate the efficiency rationale for slowing down automation. To this end, we enrich our baseline model along several dimensions that are important for the ability of workers to reallocate and smooth consumption. In particular, we allow for gross flows across occupations, uninsurable idiosyncratic earnings and mobility risks, and some forms of social insurance. We also introduce gradual automation and overlapping generations of workers (as in Section 4.6.2). Appendix B provides further details.

5.1 Firm

Production. There is a continuum of occupations of mass 1. A share ϕ are automatable ($h = A$) and a share $1 - \phi$ are not ($h = N$). Occupations use the technology

$$y_t^A = A^A (\alpha + \mu^A)^{1-\eta} \quad \text{and} \quad y_t^N = A^N (\mu^N)^{1-\eta} \quad (5.1)$$

where $\eta \in (0, 1)$ is the span of control, and $A^h > 0$ are occupation-specific productivities.²⁹ The firm's final good technology is

$$G(y_t^A, y_t^N) = \left[\phi (y_t^A)^{\frac{\nu-1}{\nu}} + (1 - \phi) (y_t^N)^{\frac{\nu-1}{\nu}} \right]^{\frac{\nu}{\nu-1}}, \quad (5.2)$$

²⁹ We normalize the relative productivity of automation to 1. This is without loss of generality since only the ratio between this productivity and the cost automation q_t is relevant.

where $\nu < 1$ is the elasticity of substitution.

Investment. The firm invests in automation. The law of motion of automation is

$$d\alpha_t = (x_t - \delta\alpha_t) dt, \quad (5.3)$$

where δ is the rate of depreciation, and x_t is the investment rate. The firm incurs a resource cost q_t per unit of investment x_t . As in [Guerreiro et al. \(2022\)](#), we assume that this cost falls over time $q_t = q^{\text{fin}} + \exp(-\psi t) (q^{\text{init}} - q^{\text{fin}})$, where q^{init} and q^{fin} are the initial and final costs, and $\psi > 0$ is the convergence rate. The initial cost ensures that automation starts at $\alpha_0 = 0$. The government taxes automation linearly at rate $\{\tau_t^x\}$ and rebates the proceedings to the firm.

Dividends. The firm smoothes dividends over time ([Leary and Michaely, 2011](#)) and issues debt to finance investment early on.³⁰ Dividends are given by $\Pi_t^{\text{div}} = \Pi^{\text{fin}} + \exp(-\vartheta t) (\Pi^{\text{init}} - \Pi^{\text{fin}})$ where Π^{init} and Π^{fin} are profits at the initial and final steady states. The convergence rate $\vartheta > 0$ ensures that the firm repays its debt $\int_0^{+\infty} \exp\left(-\int_0^t r_s ds\right) (\Pi_t^{\text{div}} - \Pi_t) = 0$.

5.2 Workers

There are overlapping generations of workers that are replaced at rate χ .³¹ A worker is indexed by five states: their asset holdings (a); their occupation of employment (h); their employment status (e); their permanent productivity component (ξ); and the mean-reverting component of their productivity (z). We let $\mathbf{x} \equiv (a, h, e, \xi, z)$ be the workers' states and π be its measure.

Assets and constraints. The asset structure is the same as in our baseline model.³² In

³⁰ Assuming that the firm smoothes dividends is conservative with respect to our mechanism. At short horizons, the investment cost of automation exceeds revenues and profits are negative. Disbursing negative dividends to workers would make them more likely to become borrowing constrained. This would strengthen the efficiency rationale for taxing automation and produce larger welfare gains. Our specification ensures that dividends do not fall; firms seem to be reluctant to cut them in practice ([Leary and Michaely, 2011](#)).

³¹ We introduce overlapping generations because young cohorts account for a substantial share of labor reallocation across occupations ([Adão et al., 2020](#)).

³² The mutual fund now rebates dividends Π_t^{div} to workers, instead of profits.

addition, workers have access to annuities which allow them to self-insure against survival risk. Financial markets are otherwise incomplete: workers cannot trade contingent securities against the risk that their occupation becomes automated, against the risk that they are not able to relocate, against unemployment risk, or against idiosyncratic productivity risk. Workers now face the budget constraint

$$da_t(\mathbf{x}) = [\mathcal{Y}_t^{\text{net}}(\mathbf{x}) + (r_t + \chi) a_t(\mathbf{x}) - c_t(\mathbf{x})] dt \quad (5.4)$$

where $\mathcal{Y}_t^{\text{net}}(\mathbf{x})$ denotes net income and r_t is the return on bonds. Workers still face the borrowing constraint $a_t(\mathbf{x}) \geq \underline{a}$. They hold $a^{\text{birth}}(\mathbf{x}) = 0$ assets at birth.

Occupational choice. Workers choose their first occupation of employment at birth. They supply labor and are given the opportunity to move between occupations with intensity λ . Moreover, workers are subject to linearly additive taste shocks when choosing between occupations. These taste shocks are distributed according to an Extreme Value Type-I distribution with mean 0 and scale parameter $\gamma > 0$, as is standard in the literature (Artuç et al., 2010). In particular, workers choose a *non-automated* occupation with hazard

$$\mathcal{S}_t(\mathbf{x}) = \frac{(1 - \phi) \exp\left(\frac{V_t^N(\mathbf{x}'(N; \mathbf{x}))}{\gamma}\right)}{\sum_{h'} \phi^{h'} \exp\left(\frac{V_t^{h'}(\mathbf{x}'(h'; \mathbf{x}))}{\gamma}\right)}, \quad (5.5)$$

where $V_t^h(\cdot)$ denotes the continuation value associated to automated ($h = A$) and non-automated ($h = N$) occupations, $\mathbf{x}'$ captures the worker's new state after they choose an occupation, and the parameter γ governs the elasticity of labor supply. Workers who reallocate go through unemployment spells which they exit at rate κ . Upon entering their new occupation, workers experience a permanent productivity loss θ . They experience this loss only the first time they reallocate.

Income. Employed workers ($e = E$) earn gross labor income

$$\mathcal{Y}_t^{\text{labor}}(\mathbf{x}) = \xi \exp(z) w_t^h, \quad (5.6)$$

with the productivity consisting of a permanent component (ξ) and a mean-reverting

component (z). The permanent component switches from 1 to $1 - \theta$ the first time a worker switches occupations. The mean-reverting component of productivity evolves as

$$dz_t = -\rho_z z_t dt + \sigma_z dW_t \quad (5.7)$$

with persistence $\rho_z^{-1} > 0$ and volatility $\sigma_z > 0$. The employment status switches to $e_t = U$ upon reallocation and reverts to $e_t = E$ upon exiting unemployment. All workers are born with $e_t = E$. As in [Krueger et al. \(2016\)](#), unemployed workers ($e = U$) receive unemployment benefits that are proportional to the gross labor income they would have earned in their previous occupation. The replacement rate is $b \in [0, 1]$. Workers claim dividends in proportion to their idiosyncratic (mean-reverting) productivity, as in [Auclert et al. \(2018\)](#).³³ Workers net income is

$$\mathcal{Y}_t^{\text{net}}(\mathbf{x}) = \mathcal{T}_t \left(\mathcal{Y}_t^{\text{labor}}(\mathbf{x}) + \exp(z) \Pi_t^{\text{div}} \right), \quad (5.8)$$

where $\mathcal{T}_t(y) = \psi_0 y^{1-\psi_1}$ captures non-linear taxation ([Heathcote et al., 2017](#)).

5.3 Policy and Equilibrium

The government's flow budget constraint is

$$dB_t = (T_t + r_t B_t - U_t - g_t) dt \quad (5.9)$$

where B_t is the government's asset holdings, T_t is total tax revenues, U_t is total unemployment benefits, and g_t is government spending. The resource constraint is now

$$\int c_t(\mathbf{x}) d\pi_t + g_t + q_t x_t = G(y_t^A, y_t^N) \quad (5.10)$$

The wages are still given by (3.1). A competitive equilibrium is defined as before.

³³ This assumption implies that workers claim labor and profit income in proportion to their idiosyncratic (mean-reverting) productivity. It is the most neutral possible, as it ensures that the government has no incentives to tax (or subsidize) automation to reduce workers' income risk.

6 Quantitative Evaluation

We now use the model to evaluate the importance of our mechanism and perform policy experiments. Section 6.1 discusses the calibration. Section 6.2 describes the laissez-faire transition. Section 6.3 discusses policy interventions. Finally, Appendices B–C provide details about our numerical implementation.

6.1 Calibration

We parameterize the model using a mix of external and internal calibration. We interpret our initial stationary equilibrium (before automation) as the year 1980. Table 6.1 shows the parameterization.

External calibration. External parameters are borrowed from the literature. The elasticity of substitution across occupations ν is 0.9 (Goos et al., 2014). The span of control parameter η is 0.15 (Basu and Fernald, 1997; Atkeson and Kehoe, 2007). The depreciation rate δ is 10%, as in Graetz and Michaels (2018). We choose an inverse elasticity of intertemporal substitution $\sigma \rightarrow 1$ as in Guerreiro et al. (2022). We set the replacement rate χ to obtain an average active lifespan of 50 years (Nuño and Moll, 2018). We pick the unemployment exit hazard parameter κ to match the average unemployment duration in the US, as measured by Alvarez and Shimer (2011). The productivity loss θ when moving between occupations is set to match the earnings losses in Kambourov and Manovskii (2009). As in Auclert et al. (2018), we rule out borrowing $\underline{a} = 0$. We use the annual income process estimated by Floeden and Lindé (2001) using PSID data and choose the persistence ρ_z^{-1} and volatility σ_z in our continuous time model accordingly. The replacement rate when unemployed b is 0.4, following Ganong et al. (2020). Government spending relative to consumption g_t/C_t is 50% at the initial steady state. The progressivity of the tax schedule ψ_1 is 0.181, as in Heathcote et al. (2017). We choose the intercept of the tax schedule ψ_0 so that the government can finance $g_t/C_t = 0.5$ at the initial steady state. Finally, the ratio of liquidity to GDP $-B_t/Y_t$ is 0.5 at the initial and final steady states (Survey of Consumer Finances, 1980).³⁴ During the transition, the

³⁴ We obtain this value by adding up checkable deposits, time and savings deposits, and money market funds share (Table B.100, lines 11-13, year 1980) and divide by nominal GDP. This amount of liquidity is almost twice as high as the one in Kaplan et al. (2018), which is conservative with

government adjusts liquidity B_t (and government spending g_t accordingly) so that the interest rate converges exponentially to its long-run level.³⁵ The half-life is the same as the one we target internally for the wage gap across occupations (15 years).

Internal calibration. We calibrate eight parameters internally: the discount rate (ρ); the mobility hazard (λ); the scale parameter (γ); the two occupation-specific productivities (A^h); the share of automated occupations (ϕ); the final cost of investment (q^{fin}); and its convergence rate (ψ). We pick these to jointly match eight moments. The discount rate targets an annual real interest rate of 2 percent. The mobility hazard λ targets an occupational mobility rate of 10% per year at the initial steady state, which roughly corresponds to the US level in 1980 in [Kambourov and Manovskii \(2008\)](#). The scale parameter γ targets an elasticity of labor supply of 1 for the stock of workers (i.e., all generations) which lies between the estimates of [Wiswall and Zafar \(2015\)](#) and [Hsieh et al. \(2019\)](#).³⁶ The occupations' productivities $\{A^h\}$ are such that output is 1 and wages are identical across occupations at the initial steady state. The mass of automated occupations ϕ targets an employment share of 56% in routine occupations in 1980 ([Bharadwaj and Dvorkin, 2019](#)). We choose the final cost of investment q^{fin} to obtain a log wage gap of 0.45 between occupations at the final steady state ([Cortes, 2016](#)). Similarly, the convergence rate ψ targets a half-life of 15 years for the wage gap ([Cortes, 2016](#)).

Untargeted moments. The model matches well several untargeted moments (see Appendix D for details). First, the share of hand-to-mouth workers is roughly 23% at the initial steady state, which lies between the estimates of [Kaplan et al. \(2014\)](#) and [Aguiar et al. \(2020\)](#). Second, the share of routine employment 40 years into the transition (year 2020) is 38% compared to roughly 41% in the data ([Bharadwaj and Dvorkin, 2019](#)). Third, we obtain that 67% of output in occupation $h = A$ is produced by automation at the final steady state. For comparison, the [McKinsey](#)

respect to our mechanism.

³⁵ The interest rate rises slowly from 2% to 2.25%. The interest rate and income taxes are essentially constant in equilibrium and do not change when taxing automation. Therefore, the welfare gains from taxing automation in Table 6.2 are not driven by redistribution or improvements in risk-sharing induced by changes in interest rate or taxes, which are unrelated to the mechanism of interest.

³⁶ We compute this elasticity in our model by simulating a 10% wage increase in one of the occupations and leaving the other one unchanged.

Table 6.1: Calibration

Parameter	Description	Calibration	Target / Source
<i>Workers</i>			
ρ	Discount rate	0.040	2% real interest rate
σ	EIS (inverse)	1	-
χ	Death rate	1/50	Average working life of 50 years
$\underline{a}$	Borrowing limit	0	Auclert et al. (2018)
<i>Technology</i>			
A^A, A^N	Productivities	(0.719, 1.710)	Initial output (1) and symmetric wages
$1 - \eta$	Initial labor share	0.85	Span of control (Atkeson and Kehoe, 2007)
δ	Depreciation rate	0.1	Graetz and Michaels (2018)
ϕ	Share of automated occupations	0.537	Routine occs. employment share in 1980
q^{fin}	Final cost of investment	5.621	Final wage gap
ψ	Convergence rate of cost	0.054	Half-life of wage gap
ν	Elasticity of subst. across occs.	0.9	Goos et al. (2014)
<i>Mobility frictions</i>			
λ	Mobility hazard	0.364	Occupational mobility rate in 1980
$1/\kappa$	Average unemployment duration	1/3.2	Alvarez and Shimer (2011)
θ	Productivity loss from relocation	0.18	Kambourov and Manovskii (2009)
γ	Scale parameter	0.0360	Elasticity of labor supply
<i>Government</i>			
ψ_0	Intercept of tax schedule	0.661	Gvt spending g/C (BEA)
ψ_1	Elasticity of tax schedule	0.181	Heathcote et al. (2017)
$-B/Y$	Liquidity / GDP	0.5	Liquid assets / GDP (SCF)
<i>Income process</i>			
ρ_z	Mean reversion	0.09	Floden and Lindé (2001)
σ_z	Volatility	0.205	Floden and Lindé (2001)
b	Replacement rate	0.4	Ganong et al. (2020)

(2017) report finds that roughly 70% of output previously produced by labor could be automated in occupations most susceptible to automation (making up for 51% of initial employment, compared to 56% in our model). Fourth, the (partial equilibrium) effects of automation on employment in our model are comparable to the firm-level estimates in Bonfiglioli et al. (2022). They find that adopting automation changes employment by -0.54 log-points at the firm level, compared to -0.66 in our model (see Appendix D for details). Finally, consumption increases by only 5.3% over the first 40 years. This is consistent with the view that automation delivers relatively small TFP gains (Acemoglu and Restrepo, 2019).

6.2 Laissez-Faire

We start by simulating the laissez-faire transition. The economy is initially at its steady state with no automation ($\alpha_0 = 0$). In period $t = 0$, the cost of investing in automation starts to fall and the economy converges gradually to its new steady state with positive automation.


Figure 6.1 shows this transition (solid lines). The rise in automation displaces workers and reallocates labor away from automated occupations. Despite this reallocation, wages decline gradually in automated occupations (red line) but increase in non-automated occupations (blue line) since the two occupations are complements. The wage gap widens to 0.45 with a half-life of 15 years (both are targeted moments). Finally, automated workers consume less and have steeper consumption profiles — their intertemporal MRS is lower. Indeed, the income of automated workers falls initially and (partially) recovers over time, as they are able to reallocate to non-automated occupations (Appendix E). This creates a motive for them to borrow against future incomes, and they are more likely to become borrowing constrained.

6.3 Second Best and Welfare

We now solve for the optimal policy and quantify welfare gains. The government maximizes

$$\mathcal{W}(\eta) \equiv \int_{-\infty}^{+\infty} \int \eta_t(\mathbf{x}) V_t^{\text{birth}}(\mathbf{x}) d\pi_t(\mathbf{x}) dt \quad (6.1)$$

Figure 6.1: Allocations



Notes: Solid curves correspond to the laissez-faire and dashed curves to the second best. ‘Wages’ are the marginal products of labor in the two occupations, and ‘Consumption’ is the average consumption by workers initially employed in a given occupation. Red and blue curves are used to denote automated and non-automated occupations / workers, respectively. Wages and consumptions are normalized by their initial steady state levels.

where $V_t^{\text{birth}}(\mathbf{x})$ is the value of a worker with state $\mathbf{x}$ born in period t , and $\eta_t(\mathbf{x})$ are Pareto weights. The government maximizes this objective by choosing taxes on investment $\{\tau_t^x\}$ along the transition. The government uses *efficiency* weights $\eta_t(\mathbf{x})$ which are inversely related to workers’ marginal utility at birth. These weights are the ones we described in our baseline model (Section 4.5) and ensure that the government has no incentive to redistribute resources. Computational details are provided in Appendix B.3.

Figure 6.1 illustrates the effects of the second best intervention (dashed lines). The optimal policy slows down automation so as to increase its half-life from 16 years at the laissez-faire to 22 years at the optimum. The speed of automation is especially slower over the first decade of the transition. There is less labor reallo-

Table 6.2: Welfare Gains ΔW from Second Best Interventions

	Benchmark	Less liquidity	Less reallocation	More complements
Automated	0.80%	0.91%	0.93%	0.78%
Non-autom.	−0.19%	−0.22%	−0.35%	−0.21%
New gener.	−0.08%	−0.11%	−0.10%	−0.08%
Total	0.20%	0.24%	0.20%	0.19%

Note: ‘Benchmark’ corresponds to the gains from the optimal taxation of automation under the calibration described in Section 6.1. ‘Less liquidity’ and ‘Less reallocation’ denote alternative calibrations where we target a ratio of liquidity to GDP of 0.35 (instead of 0.5) and a separation rate of 7.2% (instead of 10%), respectively. ‘More complements’ denotes an alternative calibration where the elasticity of substitution across occupations is 0.76 (instead of 0.9).

cation over this period, and the wage gap opens up more slowly. As anticipated in Remark 2, the optimal policy raises the income of automated workers early on during the transition when they value it more.

Table 6.2 reports the welfare gains (in consumption equivalent terms) from the second best intervention. The first column corresponds to our benchmark calibration. Automated workers benefit from slower automation (0.80%). Non-automated workers are hurt by the intervention (−0.19%). Taxing automation goes a long way in improving the welfare of automated workers. At the second best, they are only worse off by −0.60% relative to non-automated workers, compared to −1.58% at the laissez-faire. The intervention lowers slightly the welfare of new generations (−0.08%) since it reduces the value of the firm and hence dividends. Overall, the policy raises social welfare by 0.20%.³⁷ Figure F.1 in Appendix F plots the tax on investments $\{\tau_t^x\}$ that implement this second best. The optimal tax starts at roughly 5%, raises progressively to 7% over a decade, and then gradually declines to zero in the long-run.

Robustness checks. We then consider additional calibrations to assess the robustness of our results. First, we consider a narrower definition of liquid assets. We

³⁷ These figures are comparable to the ones found in the heterogeneous agent literature on optimal taxation (e.g., Heathcote et al., 2017) or the literature on the taxation of automation (e.g., Guerreiro et al., 2022).

now target a ratio $-B_t/Y_t$ of 35% (instead of 50%).³⁸ All other parameters are recalibrated to match the moments described in Section 6.1.³⁹ Automated workers are more likely to become borrowing constrained, and their consumption profiles are steeper. Therefore, they benefit more from slowing down automation and the total welfare gains increase. Second, we recognize that occupational mobility has decreased in recent decades. We thus target an occupational mobility rate of 7.2% (instead of 10%) following Moscarini and Vella (2002). This alternative calibration lowers the consumption of automated workers (in levels) compared to our benchmark as they reallocate less, but does not affect meaningfully the slope of their consumption profiles. Accordingly, automated workers benefit more from the intervention but the total welfare gains are mostly unchanged. Third, there is some uncertainty in the literature about the elasticity of substitution between occupations. We thus decrease it to 0.76 (instead of 0.9) based on Gregory et al. (2021) and find very similar welfare gains to our benchmark. Finally, we perform two more robustness checks. We increase the depreciation rate to 20% (instead of 10%) to capture the fact that some forms of automation like artificial intelligence software could depreciate faster than others like robots. We also decrease the elasticity of labor supply across occupations to 0.4 (instead of 1) following Wiswall and Zafar (2015). We obtain total welfare gains of respectively 0.19% and 0.29% (not shown in the table).

Wage supplements. Government transfers that target automated workers could in principle be an effective tool to respond to automation. In particular, the government could provide wage supplements to automated workers — similar to Trade Adjustment Assistance for Workers (TAA) in the US. This intervention would be financed by taxing non-automated workers — a negative wage supplement. We compute the wage supplements (along the transition) that would make workers indifferent between these supplements and the tax on automation. In present discounted terms, the government would need to give \$19,126 to the average auto-

³⁸ We obtain this value by subtracting consumer credit and other loans excluding mortgages (Table B.100, lines 34 and 36-37, year 1980) from our previous measure of liquidity (footnote 34).

³⁹ In particular, matching the wage dynamics in Cortes (2016) is important for the welfare gains. For instance, we have found in numerical experiments that the welfare gains are larger when the wage gap opens up faster than it did in the data. The results are robust to changes in parameterizations when matching the same wage dynamics.

ated worker, and would tax \$4,622 from the average non-automated worker.⁴⁰ Assuming a workforce of 107 million in 1980, the wage supplements to automated workers would cost roughly \$1.1 trillion and leave a fiscal deficit (after taxing non-automated workers) of roughly \$924 billion. For comparison, the US Congress budgeted \$551 million for the TAA program in 2022, or \$13.4 billion in present value in our model. These figures show that slowing down automation delivers welfare gains that would be costly to replicate with wage supplements alone.

7 Conclusion

We presented two novel results in economies where workers displaced by automation face reallocation and borrowing frictions. First, automation is inefficient when these frictions are sufficiently severe. Firms fail to internalize that automation depresses the income of automated workers early on, precisely when they become borrowing constrained. Second, the government finds it optimal to slow down automation on efficiency grounds, even when it does not value equity. Quantitatively, slowing down automation achieves meaningful welfare gains.

To derive sharp results and clarify the mechanisms at play, our model necessarily abstracted from many features. In particular, tax codes often subsidize capital and R&D expenditures on the grounds that firms face credit constraints or that there are externalities involved. Thus, our results do not necessarily imply that automation technologies ought to be taxed *on net*, as is the case for autonomous vehicles used by transportation companies in Nevada. Instead, they imply that subsidies on investment in automation should be lowered *temporarily* while the economy adjusts and displaced workers reallocate, which is similar to the lower tax credits for automation in South Korea.

Our quantitative model points to two directions for future work. First, we found that the optimal policy depends on how steep the consumption profiles of workers displaced by automation are. It would be interesting to measure these profiles in the data and use them to discipline future quantitative exercises. Second, the quantitative model is rich enough to tackle other optimal policy questions where workers are displaced and might be borrowing constrained, such as how fast governments should implement trade liberalizations or carbon taxation.

⁴⁰ Average earnings are \$65k at the initial steady state.

References

- ACEMOGLU, D., A. MANERA, AND P. RESTREPO (2020): "Does the US Tax Code Favor Automation?" *Brookings Papers on Economic Activity*.
- ACEMOGLU, D. AND P. RESTREPO (2018): "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 108, 1488–1542.
- (2019): "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 33, 3–30.
- (2022): "Tasks, Automation, and the Rise in US Wage Inequality," *Forthcoming, Econometrica*.
- ADÃO, R., M. BERAJA, AND N. PANDALAI-NAYAR (2020): "Technological Transitions with Skill Heterogeneity Across Generations," Tech. rep., National Bureau of Economic Research.
- AGHION, P. AND O. J. BLANCHARD (1994): "On the Speed of Transition in Central Europe," in *NBER Macroeconomics Annual 1994, Volume 9*, National Bureau of Economic Research, Inc, NBER Chapters, 283–330.
- AGUIAR, M. A., M. AMADOR, AND C. ARELLANO (2021): "Micro Risks and Pareto Improving Policies with Low Interest Rates," NBER Working Paper 28996.
- AGUIAR, M. A., M. BILS, AND C. BOAR (2020): "Who Are the Hand-to-Mouth?" NBER Working Papers 26643, National Bureau of Economic Research, Inc.
- AIYAGARI, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 109, 659–684.
- (1995): "Optimal Capital Income Taxation with Incomplete Markets, Borrowing Constraints, and Constant Discounting," *Journal of Political Economy*, 103, 1158–1175.
- AIYAGARI, S. R. AND E. R. MCGRATTAN (1998): "The Optimum Quantity of Debt," *Journal of Monetary Economics*, 42, 447–469.
- ALVAREZ, F. AND R. SHIMER (2011): "Search and Rest Unemployment," *Econometrica*, 79, 75–122.
- ARTUÇ, E., S. CHAUDHURI, AND J. MCLAREN (2010): "Trade Shocks and Labor Adjustment: A Structural Empirical Approach," *American Economic Review*, 100, 1008–1045.

- ATKESON, A. AND P. J. KEHOE (2007): "Modeling the Transition to a New Economy: Lessons from Two Technological Revolutions," *American Economic Review*, 97, 64–88.
- AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2018): "The Intertemporal Keynesian Cross," NBER Working Paper 25020.
- BASU, S. AND J. G. FERNALD (1997): "Returns to Scale in U.S. Production: Estimates and Implications," *Journal of Political Economy*, 105, 249–283.
- BERAJA, M. AND N. ZORZI (2022): "Inefficient Automation," Working Paper 30154, National Bureau of Economic Research.
- BHANDARI, A., D. EVANS, M. GOLOSOV, AND T. SARGENT (2021): "Efficiency, Insurance, and Redistribution Effects of Government Policies," *Working paper*.
- BHARADWAJ, A. AND M. A. DVORKIN (2019): "The Rise of Automation: How Robots May Impact the U.S. Labor Market," Regional Economist Second Quarter 2019, Federal Reserve Bank of St. Louis.
- BONFIGLIOLI, A., R. CRINO, H. FADINGER, AND G. GANCIA (2022): "Robot Imports and Firm-Level Outcomes," *Mimeo*.
- BRAXTON, J. C. AND B. TASKA (2023): "Technological Change and the Consequences of Job Loss," *American Economic Review*, 113, 279–316.
- CABALLERO, R. J. AND M. L. HAMMOUR (1996): "On the Timing and Efficiency of Creative Destruction," *The Quarterly Journal of Economics*, 111, 805–852.
- CALVO, G. A. AND M. OBSTFELD (1988): "Optimal Time-Consistent Fiscal Policy with Finite Lifetimes," *Econometrica*, 56, 411–432.
- CARD, D., J. KLUVE, AND A. WEBER (2018): "What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations," *Journal of the European Economic Association*, 16, 894–931.
- CARROLL, C. D. (2006): "The method of endogenous gridpoints for solving dynamic stochastic optimization problems," *Economics Letters*, 91, 312–320.
- CHAMLEY, C. (1986): "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives," *Econometrica*, 54, 607–622.
- CHETTY, R. (2008): "Moral Hazard versus Liquidity and Optimal Unemployment Insurance," *Journal of Political Economy*, 116, 173–234.
- CONESA, J. C., S. KITAO, AND D. KRUEGER (2009): "Taxing Capital? Not a Bad

- Idea after All!" *American Economic Review*, 99, 25–48.
- CORTES, G. M. (2016): "Where have the middle-wage workers gone? A study of polarization using panel data," *Journal of Labor Economics*, 34, 63–105.
- COSTINOT, A. AND I. WERNING (2022): "Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation," *Forthcoming, The Review of Economic Studies*.
- CRÉPON, B. AND G. J. VAN DEN BERG (2016): "Active Labor Market Policies," *Annual Review of Economics*, 8, 521–546.
- DÁVILA, E. AND A. KORINEK (2018): "Pecuniary Externalities in Economies with Financial Frictions," *Review of Economic Studies*, 85, 352–395.
- DÁVILA, J., J. H. HONG, P. KRUSELL, AND J.-V. RÍOS-RULL (2012): "Constrained Efficiency in the Neoclassical Growth Model With Uninsurable Idiosyncratic Shocks," *Econometrica*, 80, 2431–2467.
- DIAMOND, P. A. (1965): "National Debt in a Neoclassical Growth Model," *The American Economic Review*, 55, 1126–1150.
- FLODEN, M. AND J. LINDE (2001): "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?" *Review of Economic Dynamics*, 4, 406–437.
- GANONG, P., P. NOEL, AND J. VAVRA (2020): "US Unemployment Insurance Replacement Rates During the Pandemic," *Journal of Public Economics*, 191.
- GEANAKOPOLOS, J. AND H. M. POLEMARCHAKIS (1985): "Existence, Regularity, and Constrained Suboptimality of Competitive Allocations When the Asset Market Is Incomplete," *Cowles Foundation Discussion Papers* 764.
- GOLOSOV, M. AND A. TSYVINSKI (2006): "Designing Optimal Disability Insurance: A Case for Asset Testing," *Journal of Political Economy*, 114, 257–279.
- GOOS, M., A. MANNING, AND A. SALOMONS (2014): "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," *American Economic Review*, 104, 2509–2526.
- GRAETZ, G. AND G. MICHAELS (2018): "Robots at Work," *The Review of Economics and Statistics*, 100, 753–768.
- GREGORY, T., A. SALOMONS, AND U. ZIERAHN (2021): "Racing With or Against the Machine? Evidence on the Role of Trade in Europe," *Journal of the European Economic Association*, 20, 869–906.

- GUERREIRO, J., S. REBELO, AND P. TELES (2022): "Should Robots Be Taxed? [Skills, Tasks and Technologies: Implications for Employment and Earnings]," *Review of Economic Studies*, 89, 279–311.
- GUNER, N., R. KAYGUSUZ, AND G. VENTURA (2021): "Rethinking the Welfare State," CEPR Discussion Papers 16275, C.E.P.R. Discussion Papers.
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2017): "Optimal Tax Progressivity: An Analytical Framework," *The Quarterly Journal of Economics*, 132, 1693–1754.
- HÉMOUS, D. AND M. OLSEN (2022): "The Rise of the Machines: Automation, Horizontal Innovation, and Income Inequality," *American Economic Journal: Macroeconomics*, 14, 179–223.
- HSIEH, C.-T., E. HURST, C. I. JONES, AND P. J. KLENOW (2019): "The Allocation of Talent and U.S. Economic Growth," *Econometrica*, 87, 1439–1474.
- HUGGETT, M. (1993): "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17, 953–969.
- HUMLUM, A. (2019): "Robot Adoption and Labor Market Dynamics," *Mimeo*.
- HUMLUM, A., P. P. JØRGENSEN, AND J. R. MUNCH (2022): "Changing Tracks: Human Capital Investment after Loss of Earning Ability," .
- HUMMELS, D., J. R. MUNCH, AND C. XIANG (2018): "Offshoring and labor markets," *Journal of Economic Literature*, 56, 981–1028.
- ITSKHOKI, O. AND B. MOLL (2019): "Optimal Development Policies With Financial Frictions," *Econometrica*, 87, 139–173.
- JACOBSON, L. S., R. J. LALONDE, AND D. G. SULLIVAN (1993): "Earnings Losses of Displaced Workers," *American Economic Review*, 83, 685–709.
- (2005): "Is Retraining Displaced Workers a Good Investment?" *Economic Perspectives*, 29, 47–66.
- JAPPELLI, T. AND L. PISTAFERRI (2017): *The Economics of Consumption: Theory and Evidence*, no. 9780199383153 in OUP Catalogue, Oxford University Press.
- JUDD, K. L. (1985): "Redistributive Taxation in a Simple Perfect Foresight Model," *Journal of Public Economics*, 28, 59–83.
- KAMBOUROV, G. AND I. MANOVSKII (2008): "Rising Occupational And Industry Mobility In The United States: 1968-97," *International Economic Review*, 49, 41–79.

- (2009): “Occupational Specificity Of Human Capital,” *International Economic Review*, 50, 63–115.
- KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): “Monetary Policy According to HANK,” *American Economic Review*, 108, 697–743.
- KAPLAN, G., G. L. VIOLANTE, AND J. WEIDNER (2014): “The Wealthy Hand-to-Mouth,” *Brookings Papers on Economic Activity*, 45, 77–153.
- KORINEK, A. AND J. E. STIGLITZ (2020): “Steering Technological Progress,” Mimeo, University of Virginia.
- KOVACEV, R. (2020): “A Taxing Dilemma: Robot Taxes and the Challenges of Effective Taxation of AI, Automation and Robotics in the Fourth Industrial Revolution,” *Ohio St. Tech. LJ*, 16, 182.
- KRUEGER, D., K. MITMAN, AND F. PERRI (2016): “Macroeconomics and Household Heterogeneity,” vol. 2 of *Handbook of Macroeconomics*, chap. 0, 843–921.
- LANDAIS, C. AND J. SPINNEWIJN (2021): “The Value of Unemployment Insurance,” *The Review of Economic Studies*, 88, 3041–3085.
- LEARY, M. T. AND R. MICHAELY (2011): “Determinants of Dividend Smoothing: Empirical Evidence,” *The Review of Financial Studies*, 24, 3197–3249.
- LEE, D. AND K. I. WOLPIN (2006): “Intersectoral Labor Mobility and the Growth of the Service Sector,” *Econometrica*, 74, 1–46.
- LOHR, S. (2022): “Economists Pin More Blame on Tech for Rising Inequality,” *The New York Times*.
- MANKIW, N. G. AND S. ZELDES (1991): “The Consumption of Stockholders and Nonstockholders,” *Journal of Financial Economics*, 29, 97–112.
- MARTINEZ, J. (2018): “Automation, Growth and Factor Shares,” in *SED Meeting Papers*, Society for Economic Dynamics, vol. 736.
- MCKINSEY (2017): “A Future that Works: Automation, Employment, and Productivity,” *McKinsey Global Institute*.
- MOLL, B., L. RACHEL, AND P. RESTREPO (2021): “Uneven Growth: Automation’s Impact on Income and Wealth Inequality,” NBER Working Papers 28440.
- MOSCARINI, G. AND F. VELLA (2002): “Aggregate worker reallocation and occupational mobility in the United States: 1971-2000,” IFS Working Papers W02/18, Institute for Fiscal Studies.

- MOSCARINI, G. AND F. G. VELLA (2008): "Occupational Mobility and the Business Cycle," NBER Working Papers 13819.
- MUSSA, M. (1984): "The Adjustment Process and the Timing of Trade Liberalization," NBER Working Papers 1458, National Bureau of Economic Research, Inc.
- NEARY, J. P. (1982): "Intersectoral Capital Mobility, Wage Stickiness, and the Case for Adjustment Assistance," in *Import Competition and Response*, National Bureau of Economic Research, Inc, NBER Chapters, 39–72.
- NUÑO, G. AND B. MOLL (2018): "Social optima in economies with heterogeneous agents," *Review of Economic Dynamics*, 28, 150–180.
- PIKETTY, T. AND E. SAEZ (2013): "Chapter 7 - Optimal Labor Income Taxation," vol. 5 of *Handbook of Public Economics*, 391–474.
- POTERBA, J. M. (1988): "Are Consumers Forward Looking? Evidence from Fiscal Experiments," *American Economic Review*, 78, 413–418.
- (2003): "Employer Stock and 401(k) Plans," *American Economic Review*, 93, 398–404.
- ROUWENHORST, G. (1995): *Asset Pricing Implications of Equilibrium Business Cycle Models*, vol. Frontiers of Business Cycle Research, Princeton University Press.
- SCHOCHET, P. Z., R. D'AMICO, J. BERK, S. DOLFIN, AND N. WOZNY (2012): "Estimated Impacts for Participants in the Trade Adjustment Assistance (TAA) Program Under the 2002 Amendments," Tech. rep., Mathematica Policy Research.
- SULLIVAN, J. X. (2008): "Borrowing During Unemployment: Unsecured Debt as a Safety Net," *Journal of Human Resources*, 43, 383–412.
- THUEMMEL, U. (2018): "Optimal Taxation of Robots," CESifo Working Paper 7317.
- WERNING, I. (2015): "Incomplete Markets and Aggregate Demand," NBER Working Papers 21448, National Bureau of Economic Research, Inc.
- WISWALL, M. AND B. ZAFAR (2015): "Determinants of College Major Choice: Identification using an Information Experiment," *Review of Economic Studies*, 82, 791–824.
- YOUNG, E. R. (2010): "Solving the incomplete markets model with aggregate uncertainty using the Krusell-Smith algorithm and non-stochastic simulations," *Journal of Economic Dynamics and Control*, 34, 36–41.

Online Appendix for: Inefficient Automation

This online appendix contains the proofs and derivations of all theoretical results for the article “Inefficient Automation,” as well as a detailed description of the quantitative model and how it is solved numerically.

Any references to equations, figures, tables, assumptions, propositions, lemmas, or sections that are not preceded “A.,” “B.,” “C.,” “D.,” “E.” or “F.” refer to the main article.

A Proofs and Derivations

A.1 Proof of Lemma 1

Fix some period $T \geq 0$. Consider the decision of automated workers to reallocate, i.e. the choice of $\{m_t\}$. Using a standard variational argument, it is optimal to reallocate for all workers who are able to ($m_t = 1$) if and only if the present discounted value of their labor income is higher in the non-automated occupation

$$\int_T^{+\infty} \exp(-\rho(t-T)) u'(c_t^A) \Delta_t dt > 0, \quad (\text{A.1})$$

where

$$\Delta_t \equiv (1-\theta) [\iota(1 - \exp(-\kappa(t-T))) + 1 - \iota] w_t^N - w_t^A \quad (\text{A.2})$$

captures the marginal increase in output from reallocating an additional worker, since $w_t^h = \partial_h G^*(\mu^A, \mu^N; \alpha)$ in equilibrium. Workers do not reallocate ($m_t = 0$) if and only if the inequality (A.1) is reversed. Any $m_t \in [0, 1]$ is optimal otherwise.

In equilibrium, reallocation takes the following form. Workers reallocate until $T^{\text{LF}} \geq 0$, i.e., $m_t = 1$ for all $t \in [0, T^{\text{LF}})$, and they stop reallocating afterwards, i.e., $m_t = 0$ for all $t \geq T^{\text{LF}}$. The reason is that the wage in automated occupations w_t^A increases over time as workers leave this occupation (by decreasing returns), and the wage in non-automated occupations w_t^N decreases as workers enter this occupation.

We next show that reallocation does take place in equilibrium, i.e., $T^{\text{LF}} > 0$. It suffices to show that workers find it optimal to reallocate at $t = 0$. That is,

$$\int_0^{+\infty} (1-\theta) [\iota(1 - \exp(-\kappa t)) + 1 - \iota] \frac{\exp(-\rho t) u'(\tilde{c}_t^A) \tilde{w}_t^N}{\int_0^{+\infty} \exp(-\rho s) u'(\tilde{c}_s^A) \tilde{w}_s^A ds} dt > 1 \quad (\text{A.3})$$

where $\{\tilde{c}_t^A\}$ and $\{\tilde{w}_t^h\}$ are counterfactual sequences of consumption and wages associated with $T = 0$ and $\alpha = \alpha^{\text{LF}}$. Consumption and wages are constant over time when $T = 0$, so the inequality (A.3) holds if and only if

$$\frac{(1-\theta)(1-\iota) \partial_N G^*\left(\frac{1}{2}, \frac{1}{2}; \alpha\right)}{\partial_A G^*\left(\frac{1}{2}, \frac{1}{2}; \alpha\right)} \frac{\rho(1-\iota) + \kappa}{(1-\iota)(\rho + \kappa)} > 1, \quad (\text{A.4})$$

where $\alpha = \alpha^{\text{LF}}$. This necessarily holds by Assumption 3 when Z^* is sufficiently small since $\partial_N G^* \left(\frac{1}{2}, \frac{1}{2}; \alpha \right) > \partial_A G^* \left(\frac{1}{2}, \frac{1}{2}; \alpha \right)$ with automation $\alpha > 0$ by Assumption 1.

A.2 Proof of Lemma 2

We begin by showing that automated workers borrow and non-automated workers save in equilibrium. We then show that automated workers become borrowing constrained when borrowing and reallocation frictions are sufficiently severe, and characterize the threshold $a^*(\lambda, \kappa)$. We focus on the case $\underline{a} < 0$ since the statement is obviously true when $\underline{a} = 0$.

Assets. It suffices to prove that $da_t^N \geq da_t^A$ for any period t where $a_t^N = a_t^A$ with strict inequality in period $t = 0$. The reason is that the equilibrium is continuous in time t , so the sequence of assets of automated and non-automated would intersect before the inequality reverses. This would imply that automated workers borrow and non-automated workers save as $a_t^N + a_t^A = 0$ in equilibrium.

To derive a contradiction, suppose instead that $da_t^N < da_t^A$ when $a_t^N = a_t^A = 0$. Then, there exists some $S > t$ such that $a_S^A > 0$ and $a_S^N < 0$ but all workers are still unconstrained $a_s^h > \underline{a}$. In this case, workers' consumptions satisfy the Euler equation

$$c_s^h = c_t^h \exp \left(\frac{1}{\sigma} \left(\int_t^s (r_\tau - \rho) d\tau \right) \right) \quad (\text{A.5})$$

for all $s \in [t, S)$. Using the market clearing condition (2.13), it must also be that

$$\exp \left(\frac{1}{\sigma} \left(\int_t^s (r_\tau - \rho) d\tau \right) \right) = \frac{\frac{1}{2} (c_s^A + c_s^N)}{\frac{1}{2} (c_t^A + c_t^N)} = \frac{C_s}{C_t} \equiv \frac{G^* (\mu_s^A, \mu_s^N; \alpha)}{G^* (\mu_t^A, \mu_t^N; \alpha)}, \quad (\text{A.6})$$

for all $s \in [t, S)$. Using the budget constraint (2.7), consumption is

$$c_t^h = \frac{\int_t^S \exp \left(- \int_t^s r_\tau d\tau \right) (\hat{Y}_s^h + \Pi_s) ds + a_t^h - \exp \left(- \int_t^S r_\tau d\tau \right) a_S^h}{\int_t^S \exp \left(- \int_t^s r_\tau d\tau \right) \exp \left(\frac{1}{\sigma} \int_t^s (r_\tau - \rho) d\tau \right) ds}, \quad (\text{A.7})$$

so assets accumulate according to

$$da_t^h = \left(\hat{\mathcal{Y}}_t^h + \Pi_t - \frac{\int_t^S \exp(-\int_t^s r_\tau d\tau) (\hat{\mathcal{Y}}_s^h + \Pi_s) ds}{\int_t^S \exp(-\int_t^s r_\tau d\tau) \exp\left(\frac{1}{\sigma} \int_t^s (r_\tau - \rho) d\tau\right) ds} + \Gamma_{t,S} a_t^h - \Gamma_{t,S}^* a_S^h \right) dt \quad (\text{A.8})$$

for some $\Gamma_{t,S}, \Gamma_{t,S}^* > 0$ that depend on the sequence of interest rates. Using (A.8),

$$\frac{d(a_t^N - a_t^A)}{C_t} = \left(z_t - \frac{\int_t^S \exp(-\int_t^s r_\tau d\tau) \frac{C_s}{C_t} z_s ds}{\int_t^S \exp(-\int_t^s r_\tau d\tau) \exp\left(\frac{1}{\sigma} \int_t^s (r_\tau - \rho) d\tau\right) ds} - \Gamma_{t,S}^* \left(\frac{a_S^N - a_S^A}{C_t} \right) \right) dt \quad (\text{A.9})$$

when $a_t^N = a_t^A = 0$, with flows $z_t \equiv (\hat{\mathcal{Y}}_t^N - \hat{\mathcal{Y}}_t^A) / C_t$. Using (A.6),

$$\frac{d(a_t^N - a_t^A)}{C_t} = \left(z_t - \int_t^S \psi_{t,s} z_s ds - \Gamma_{t,S}^* \left(\frac{a_S^N - a_S^A}{C_t} \right) \right) dt, \quad (\text{A.10})$$

with weights

$$\psi_{t,s} \equiv \frac{\exp(-\rho(s-t)) \left(\frac{C_s}{C_t}\right)^{1-\sigma}}{\int_t^S \exp(-\rho(\tau-t)) \left(\frac{C_\tau}{C_t}\right)^{1-\sigma} d\tau} > 0 \quad (\text{A.11})$$

that integrate to $\int_t^S \psi_{t,s} ds = 1$. As we will establish at the end of this appendix, $\{z_s\}$ is positive and decreases over time. The reason is twofold. First, the labor income of automated workers is lower than that of non-automated workers, and the former increases over time while the latter decreases. Second, aggregate consumption grows over time too. Therefore, $z_t - \int_t^S \psi_{t,s} z_s ds > 0$. Furthermore, $a_S^N < a_S^A$ under our postulate. It follows that $d(a_t^N - a_t^A) > 0$. This contradicts our postulate that $da_t^N < da_t^A$. This establishes that $da_t^N \geq da_t^A$ when $a_t^N = a_t^A = 0$. Repeating the steps above, the inequality is strict $da_t^N > da_t^A$ after the shock in $t = 0$. This shows that automated workers borrow in equilibrium.

Threshold $a^*(\lambda, \kappa)$. Integrating the budget constraint (2.7) over time and using (A.5) gives the assets of automated workers if they were never to become borrowing constrained

$$a_t^A = \int_0^t \exp\left(\int_s^t r_\tau d\tau\right) \left[\hat{\mathcal{Y}}_s^A + \Pi_s - c_0^A \exp\left(\frac{1}{\sigma} \int_0^s (r_\tau - \rho) d\tau\right) \right] ds. \quad (\text{A.12})$$

The sequence $\{a_t^A\}$ depends on reallocation frictions (λ, κ) but not the borrowing limit $\underline{a}$. Let $a^*(\lambda, \kappa) \equiv \inf_t a_t^A$ be the lowest value attained by this sequence. If the borrowing limit is sufficiently tight that $\underline{a} > a^*(\lambda, \kappa)$, then automated workers would become borrowing constrained in equilibrium. This shows that $\underline{a} > a^*(\lambda, \kappa)$ is a sufficient condition for borrowing constraints to bind. It is also a necessary condition because, if borrowing constraints bind, then it must be that the borrowing limit $\underline{a}$ is above $\inf_t a_t^A$. Non-automated workers never become borrowing constrained since they save in equilibrium.

Finally, we show that $a^*(\lambda, \kappa) < 0$ (i.e., borrowing constraints can bind) if and only if reallocation is slow ($1/\lambda > 0$ or $1/\kappa > 0$). To prove sufficiency, note that the model is static when reallocation is instantaneous ($1/\lambda \rightarrow 0$ and $1/\kappa \rightarrow 0$). Then, all labor income and profit changes are permanent, automated workers do not borrow, and therefore $a^*(\lambda, \kappa) \equiv \inf_t a_t^A \rightarrow 0$. To prove necessity, note that automated workers borrow $a^*(\lambda, \kappa) \equiv \inf_t a_t^A < 0$ when reallocation is slow $1/\lambda > 0$ or $1/\kappa > 0$. The reason is that $\{z_s\}$ is strictly positive and strictly decreasing over time so that automated workers borrow by (A.10). In this case, there is always a (small) borrowing limit $\underline{a} > 0$ such that automated workers become borrowing constrained.

Assumption 3. We have supposed so far that the sequence $z_t \equiv (\hat{\mathcal{Y}}_t^N - \hat{\mathcal{Y}}_t^A) / C_t$ is positive and decreases over time. The fact that $z_t > 0$ follows directly from Assumption 2 and Lemma 1. That is, automation drives a wedge between the marginal productivities of labor across occupations, and reallocation stops before the wages are fully equalized. As we show below, a sufficient condition for z_t to decrease over time is that the productivity loss θ and the duration of unemployment spells $1/\kappa$ are sufficiently small that output still increases over time.

Output increases over time when

$$\partial_t G^* (\mu_t^A, \mu_t^N; \alpha) = \partial_A G^* (\cdot) \partial_t \mu_t^A + \partial_N G^* (\cdot) \partial_t \mu_t^N > 0, \quad (\text{A.13})$$

with $\partial_t \mu_t^h$ given by the effective labor supplies (3.4)–(3.5). The condition (A.13) holds in the limit where the productivity loss of reallocation and the duration of unemployment spells are sufficiently small $1 - (1 - \theta)(1 - 1/\kappa) \rightarrow 0$ since $\partial_t \mu_t^N = -\partial_t \mu_t^A > 0$ in this case and $\partial_A G^* (\cdot) > \partial_N G^* (\cdot)$. Note that μ_t^A, μ_t^N and α are continuous in $(\theta, 1/\kappa)$ at the laissez-faire. Therefore, there exists some threshold $Z^* > 0$ such that (A.13) still holds for all $(\theta, 1/\kappa)$ such that $1 - (1 - \theta)(1 - 1/\kappa) < Z^*$.

It remains to show that the sequence $\{z_t\}$ decreases over time when $1 - (1 - \theta) \times (1 - 1/\kappa) < Z^*$. It suffices to show that $\partial_t (\hat{\mathcal{Y}}_t^N - \hat{\mathcal{Y}}_t^A) < 0$, as output and consumption C_t increase over time when this condition holds. Using labor incomes (2.8) and the effective labor supplies (3.4)–(3.5),

$$\frac{1}{2} (\hat{\mathcal{Y}}_t^N - \hat{\mathcal{Y}}_t^A) = -\partial_A G^* (\cdot) \mu_t^A + (1 - \mu_t^N) \partial_N G^* (\cdot). \quad (\text{A.14})$$

Therefore,

$$\begin{aligned} \frac{1}{2} \partial_t (\hat{\mathcal{Y}}_t^N - \hat{\mathcal{Y}}_t^A) = & - \left\{ \partial_{AA}^2 G^* (\cdot) \partial_t \mu_t^A + \partial_{AN}^2 G^* (\cdot) \partial_t \mu_t^N \right\} \mu_t^A - \partial_A G^* (\cdot) \partial_t \mu_t^A \\ & + (1 - \mu_t^N) \left\{ \partial_{NA}^2 G^* (\cdot) \partial_t \mu_t^A + \partial_{NN}^2 G^* (\cdot) \partial_t \mu_t^N \right\} \\ & - \partial_N G^* (\cdot) \partial_t \mu_t^N. \end{aligned} \quad (\text{A.15})$$

And so,

$$\frac{1}{2} \partial_t (\hat{\mathcal{Y}}_t^N - \hat{\mathcal{Y}}_t^A) < - \left\{ \partial_A G^* (\cdot) \partial_t \mu_t^A + \partial_N G^* (\cdot) \partial_t \mu_t^N \right\} \quad (\text{A.16})$$

using $\partial_{AA}^2 G^* (\cdot) < 0$ and $\partial_{NN}^2 G^* (\cdot) < 0$ since G^* is neoclassical, $\partial_{AN}^2 G^* (\cdot) > 0$ by Assumption 1, and $\partial_t \mu_t^A < 0$ and $\partial_t \mu_t^N > 0$ in equilibrium. Thus, $\partial_t (\hat{\mathcal{Y}}_t^N - \hat{\mathcal{Y}}_t^A) < 0$ in the limit $1 - (1 - \theta)(1 - 1/\kappa) \rightarrow 0$ since $\partial_t \mu_t^N = -\partial_t \mu_t^A > 0$ in this case, using the fact that $\partial_A G^* (\cdot) < \partial_N G^* (\cdot)$ in equilibrium. By continuity of the equilibrium, we still have $\partial_t (\hat{\mathcal{Y}}_t^N - \hat{\mathcal{Y}}_t^A) < 0$ when $1 - (1 - \theta) \times (1 - 1/\kappa) < Z^*$ for Z^* small enough. Taken together, the inequalities (A.13) and (A.16) imply that $z_t = (\hat{\mathcal{Y}}_t^N - \hat{\mathcal{Y}}_t^A) / C_t$ decreases over time, which completes the proof.

A.3 Proof of Lemma 3

In equilibrium, there is no arbitrage between bonds and equity since workers can trade both, so $Q_t = \exp\left(-\int_0^t r_s ds\right)$. Appendix A.2 has shown that non-automated workers are on their Euler equation $\exp(-\rho t) u'(c_t^N) / u'(c_0^N) = \exp\left(-\int_0^t r_s ds\right)$. Next, we characterize the automation choice. Using a standard variational argument, a necessary condition for an interior optimum is

$$\int_0^{+\infty} \exp\left(-\int_0^t r_s ds\right) \frac{d}{d\alpha} \Pi_t(\alpha) dt = 0. \quad (\text{A.17})$$

Furthermore, the following envelope condition applies

$$\frac{d}{d\alpha} \Pi_t(\alpha) = \partial_\alpha G^*(\mu_t^A, \mu_t^N; \alpha). \quad (\text{A.18})$$

Therefore, the following condition is necessary

$$\int_0^{+\infty} \exp\left(-\int_0^t r_s ds\right) \partial_\alpha G^*(\mu_t^A, \mu_t^N; \alpha) dt = 0. \quad (\text{A.19})$$

It is also sufficient by Assumption 2. The degree of automation is unique and interior given the labor allocations across occupations (μ_t^A, μ_t^N) .

Finally, we show that Δ_t^* increases over time in equilibrium. By definition, $\Delta_t^* \equiv \partial_\alpha G^*(\cdot)$. Therefore,

$$\partial_t \Delta_t^* = \partial_{\alpha A}^2 G^*(\mu_t^A, \mu_t^N; \alpha) \times \partial_t \mu_t^A + \partial_{\alpha N}^2 G^*(\mu_t^A, \mu_t^N; \alpha) \times \partial_t \mu_t^N \quad (\text{A.20})$$

Note that the first term on the right-hand side of (A.20) is positive, using Assumption 1, the definition of $G^*(\cdot)$ in (2.3) and the fact that $\partial_t \mu_t^A < 0$. Furthermore,

$$\begin{aligned} \partial_{\alpha N}^2 G^*(\mu_t^A, \mu_t^N; \alpha) &= \partial_{y^A, y^N}^2 G\left(F(\mu_t^A; \alpha), F(\mu_t^N; 0)\right) \\ &\quad \times \partial_\alpha F(\mu_t^A; \alpha) \times \partial_\mu F(\mu_t^N; 0) > 0, \end{aligned} \quad (\text{A.21})$$

where the inequality uses Assumption 1. Therefore, the second term on the right-hand side of (A.20) is also positive since $\partial_t \mu_t^N > 0$. This shows that $\partial_t \Delta_t^* > 0$ so Δ_t^* indeed increases over time.

A.4 Proof of Proposition 1

The result consists of two parts. First, we prove that the equilibrium is generically constrained inefficient by showing that there exists a Pareto improvement. Second, we show that the Pareto improvement involves taxing automation ($\delta\alpha < 0$) when the duration of unemployment spells $1/\kappa$ is sufficiently short and the productivity loss θ is sufficiently small (Assumption 3).

Part I (generic constrained inefficiency). The changes in welfare starting from the laissez-faire are

$$\begin{aligned} \delta U^h = & \delta\alpha \times \int_0^{+\infty} \exp(-\rho t) u'(c_t^h) (\hat{c}_t^{h,*} + \bar{c}^{h,*}) dt \\ & + \delta T \times \int_0^{+\infty} \exp(-\rho t) u'(c_t^h) (\hat{c}_t^h + \bar{c}^h) dt, \end{aligned}$$

for automated and non-automated workers $h = A, N$. The time-varying effects $\hat{c}_t^{h,*}$ and the permanent effects $\bar{c}^{h,*}$ of automation are defined in (4.4). The time-varying effects $\hat{c}_t^h$ and the permanent effects $\bar{c}^h$ of reallocation are defined similarly⁴¹

$$\hat{c}_t^h \equiv \partial_T c_t^h - \bar{c}^h \quad \text{and} \quad \bar{c}^h \equiv \int_0^{+\infty} \frac{\exp(-\rho t) u'(c_t^A)}{\int_0^{+\infty} \exp(-\rho s) u'(c_s^A) ds} \partial_T c_t^h dt. \quad (\text{A.22})$$

Using the definition of labor earnings (2.8) and profits (2.5), the consumption changes for non-automated workers are⁴²

$$\partial_\alpha c_t^N = \hat{c}_t^{N,*} + \bar{c}^{N,*} = \partial_\alpha w_t^N + \Delta_t^* - \sum_h \mu_t^h \partial_\alpha w_t^h \quad (\text{A.23})$$

$$\partial_T c_t^N = \hat{c}_t^N + \bar{c}^N = \partial_T w_t^N - \sum_h \mu_t^h \partial_T w_t^h, \quad (\text{A.24})$$

where the sequences $\{\partial_\alpha w_t^h\}$ and $\{\partial_T w_t^h\}$ are the marginal effects of perturbation in α and T on wages $w_t^h \equiv \partial_h G(\cdot)$. The marginal effects on consumption for

⁴¹ The only difference with (4.4) is that the permanent effect in (A.22) is computed using the marginal utility of A workers (who make the reallocation choice) instead of N workers (who effectively make the automation choice). This ensures that the permanent effects $\bar{c}^{h,*}$ and $\bar{c}^h$ are purely distributional $\bar{c}^{A,(\cdot)} + \bar{c}^{N,(\cdot)} = 0$ and hence not a source of inefficiency.

⁴² Expression (A.24) already use the fact that the firm chooses labor demand optimally.

automated workers ($\partial_\alpha c_t^A$ and $\partial_T c_t^A$) follow from the aggregate resource constraint

$$\frac{1}{2} \left(\partial_\alpha c_t^A + \partial_\alpha c_t^N \right) = \Delta_t^* \quad \text{and} \quad \frac{1}{2} \left(\partial_T c_t^A + \partial_T c_t^N \right) = \Delta_t. \quad (\text{A.25})$$

The permanent effects $\bar{c}^{h,*}$ and $\bar{c}^h$ are purely distributional $\bar{c}^{A,(\star)} + \bar{c}^{N,(\star)} = 0$.

There exists a perturbation $(\delta\alpha, \delta T)$ with $\delta\alpha \neq 0$ that results in $\delta U^N = 0$ and $\delta U^A > 0$ if and only if

$$\begin{aligned} & \int_0^{+\infty} \exp(-\rho t) u'(c_t^A) \left(\hat{c}_t^{A,*} + \bar{c}^{A,*} \right) dt \\ & \neq \underbrace{\frac{\int_0^{+\infty} \exp(-\rho t) u'(c_t^N) \left(\hat{c}_t^{N,*} + \bar{c}^{N,*} \right) dt}{\int_0^{+\infty} \exp(-\rho t) u'(c_t^N) \left(\hat{c}_t^N + \bar{c}^N \right) dt}}_{-\delta T / \delta\alpha \text{ that leaves } N \text{ worker indifferent}} \int_0^{+\infty} \exp(-\rho t) u'(c_t^A) \left(\hat{c}_t^A + \bar{c}^A \right) dt. \end{aligned} \quad (\text{A.26})$$

Equivalently,

$$\begin{aligned} & \int_0^{+\infty} \exp(-\rho t) u'(c_t^A) \left(\hat{c}_t^{A,*} + \bar{c}^{A,*} \right) dt \\ & \neq \underbrace{\frac{\int_0^{+\infty} \exp(-\rho t) u'(c_t^A) \left(\hat{c}_t^A + \bar{c}^A \right) dt}{\int_0^{+\infty} \exp(-\rho t) u'(c_t^N) \left(\hat{c}_t^N + \bar{c}^N \right) dt}}_{\equiv \Omega} \int_0^{+\infty} \exp(-\rho t) u'(c_t^N) \left(\hat{c}_t^{N,*} + \bar{c}^{N,*} \right) dt. \end{aligned} \quad (\text{A.27})$$

We now show that (A.27) holds with inequality *generically*. Suppose that the expression does hold with equality. Then, there exists a perturbation of the production function $G^{\star'} = \mathcal{G}(G^*, \epsilon)$ (with $\mathcal{G}(G^*, \epsilon) \rightarrow G^*$ uniformly as $\epsilon \rightarrow 0$) and a threshold $\bar{\epsilon} > 0$ such that the expression does not hold with equality in this alternative economy, for all $0 < \epsilon \leq \bar{\epsilon}$. One such perturbation is

$$\mathcal{G}(G^*, \epsilon) = G^* + \epsilon g(\mu^A, \mu^N; \alpha) \quad (\text{A.28})$$

with

$$g(\mu_t^A, \mu_t^N; \alpha) \equiv \left\{ \mu_t^A - z \right\} \left(\alpha - \alpha^{\text{LF}} \right), \quad (\text{A.29})$$

where $z > 0$ is chosen so that

$$\int_0^{+\infty} \exp(-\rho t) u'(c_t^N) \left(\mu_t^A - z \right) dt = 0. \quad (\text{A.30})$$

One can easily verify that all equilibrium conditions (Lemmas 1 and 3 and the resource constraint) are still satisfied after a perturbation $\epsilon > 0$ when evaluated at the laissez-faire. Therefore, the laissez-faire allocation is unchanged. Moreover, this perturbation ensures that (A.27) holds with inequality. To see that, note that Ω in (A.27) is unchanged after the perturbation $\epsilon > 0$. The reason is twofold. First, the terms $\hat{c}_t^N + \bar{c}^N$ are unaffected since (A.24) depends on the second order derivatives of $G^*(\cdot)$ with respect to labor (μ^A, μ^N) , while the perturbation (A.28)–(A.29) is linear in these variables. Second, the numerator is also unchanged since $\hat{c}_t^A + \bar{c}^A$ is equal to $\partial_T c_t^A$ given by (A.25).

Regarding the other terms on the left-hand and right-hand sides of (A.27), they change differently after the perturbation. To see that, let

$$\Gamma^{h,*} \equiv \frac{d}{d\epsilon} \int_0^{+\infty} \exp(-\rho t) u'(c_t^h) (\hat{c}_t^{h,*} + \bar{c}^{h,*}) dt. \quad (\text{A.31})$$

We have

$$\Gamma^{A,*} = \int_0^{+\infty} \exp(-\rho t) u'(c_t^A) \{2\mu_t^A - z\} dt > 0 \quad (\text{A.32})$$

$$\Gamma^{N,*} = - \int_0^{+\infty} \exp(-\rho t) u'(c_t^N) z dt < 0, \quad (\text{A.33})$$

using (A.23) and (A.25). To see why the inequality (A.32) holds, first note that

$$\Gamma^{A,*} > \int_0^{+\infty} \exp(-\rho t) u'(c_t^A) \{\mu_t^A - z\} dt. \quad (\text{A.34})$$

Then, let

$$\omega_t^h \equiv \frac{\exp(-\rho t) u'(c_t^h)}{\int_0^{+\infty} \exp(-\rho s) u'(c_s^h) ds} \quad (\text{A.35})$$

for each $h = A, N$. Note that the sequence $\{\omega_t^A - \omega_t^N\}$ integrates to zero and decreases over time. The reason is that the labor income (and thus consumption) of automated workers $\hat{Y}_t^A$ grows faster over time than that of non-automated workers $\hat{Y}_t^N$ (Appendix A.2). Therefore,

$$\int_0^{+\infty} \omega_t^A \{\mu_t^A - z\} dt > \int_0^{+\infty} \omega_t^N \{\mu_t^A - z\} dt = 0, \quad (\text{A.36})$$

since μ_t^A decreases over time by (3.4) and using (A.30). This shows why the inequality (A.32) holds.

Taken together, the previous steps show that (A.26) holds with inequality for virtually any economy, so that there exists a perturbation $(\delta\alpha, \delta T)$ that improves the welfare of automated workers $\delta U^A > 0$ and leaves non-automated workers indifferent $\delta U^N = 0$. That is, the equilibrium is generically constrained inefficient.

Part II (taxing automation). We now prove that the Pareto improvement requires taxing automation. The variation $(\delta\alpha, \delta T)$ with $\delta\alpha < 0$ results in $\delta U^A > 0$ and $\delta U^N = 0$ if and only if the left-hand side of inequality (A.27) is strictly less than the right-hand side

$$\begin{aligned} & \int_0^{+\infty} \exp(-\rho t) u'(c_t^A) (\hat{c}_t^{A,*} + \bar{c}^{A,*}) dt \\ & < -\frac{\delta T}{\delta\alpha} \int_0^{+\infty} \exp(-\rho t) u'(c_t^A) (\hat{c}_t^A + \bar{c}^A) dt \end{aligned} \quad (\text{A.37})$$

for automated workers (using the fact that $\delta\alpha < 0$), and

$$\begin{aligned} & \int_0^{+\infty} \exp(-\rho t) u'(c_t^N) (\hat{c}_t^{N,*} + \bar{c}^{N,*}) dt \\ & = -\frac{\delta T}{\delta\alpha} \int_0^{+\infty} \exp(-\rho t) u'(c_t^N) (\hat{c}_t^N + \bar{c}^N) dt \end{aligned} \quad (\text{A.38})$$

for non-automated workers.

To gain some intuition regarding the Pareto improvement, it is useful to rewrite the inequality (A.37) as

$$\int_0^{+\infty} \left\{ \omega_t^A - \omega_t^N \right\} \times \left\{ \hat{c}_t^{A,*} - \frac{\delta T}{\delta\alpha} \times \hat{c}_t^N \right\} dt < 0, \quad (\text{A.39})$$

where we have used (A.37)–(A.38), the definition of ω_t^h in (A.35), the definitions of the time-varying effects $\hat{c}_t^{h,(\star)}$ in (4.4) and (A.22), the fact that the permanent effects are distributional $\bar{c}^{A,(\star)} + \bar{c}^{N,(\star)} = 0$, and the fact that the ω_t^h terms integrate to 1. This inequality shows that a Pareto improvement exists *only if* borrowing constraints bind (since $\omega_t^A = \omega_t^N$ otherwise) *and* the intervention changes the timing of incomes (and hence consumptions).

Next, we prove that (A.37) indeed holds when $\delta T/\delta\alpha$ is given by (A.38) and,

therefore, the Pareto improvement requires taxing automation. Substituting $\delta T / \delta \alpha$ from (A.38) in (A.37), we can write the inequality as

$$\int_0^{+\infty} \exp(-\rho t) \left\{ \frac{u'(c_t^A)}{u'(c_T^A)} - (1 + \tilde{\Omega}) \frac{u'(c_t^N)}{u'(c_T^N)} \right\} \partial_\alpha c_t^A dt < 0 \quad (\text{A.40})$$

with

$$\tilde{\Omega} \equiv - \frac{u'(c_T^N)}{u'(c_T^A)} \frac{\int_T^{+\infty} \exp(-\rho t) u'(c_t^A) \partial_T c_t^A dt}{\int_T^{+\infty} \exp(-\rho t) u'(c_t^N) \partial_T c_t^N dt} - 1, \quad (\text{A.41})$$

where we have used the resource constraint (A.25), Lemma 3, and the fact that $\partial_T c_t^h = 0$ for all $t < T$ with $T \equiv T^{\text{LF}}$.

Suppose that $\tilde{\Omega}$ is sufficiently small. This is the case by Assumption 3, as we discuss at the end of this proof. Then, taxing automation $\delta \alpha < 0$ Pareto improves upon the laissez-faire if and only if

$$\int_0^{+\infty} \exp(-\rho t) \left\{ \frac{u'(c_t^A)}{u'(c_T^A)} - \frac{u'(c_t^N)}{u'(c_T^N)} \right\} \partial_\alpha c_t^A dt < 0, \quad (\text{A.42})$$

Using the resource constraint (A.25), we can write this inequality as

$$\begin{aligned} \int_0^T \exp(-\rho t) \left\{ \frac{u'(c_t^A)}{u'(c_T^A)} - \frac{u'(c_t^N)}{u'(c_T^N)} \right\} \times \frac{1}{2} (\partial_\alpha c_t^A - \partial_\alpha c_t^N) dt + \tilde{\Psi} \\ + \int_0^{+\infty} \exp(-\rho t) \left\{ \frac{u'(c_t^A)}{u'(c_T^A)} - \frac{u'(c_t^N)}{u'(c_T^N)} \right\} \times \Delta_t^* dt < 0, \end{aligned} \quad (\text{A.43})$$

where the term

$$\tilde{\Psi} \equiv \int_T^{+\infty} \exp(-\rho t) \left\{ \frac{u'(c_t^A)}{u'(c_T^A)} - \frac{u'(c_t^N)}{u'(c_T^N)} \right\} \times \frac{1}{2} (\partial_\alpha c_t^A - \partial_\alpha c_t^N) dt \quad (\text{A.44})$$

is sufficiently small by Assumption 3, as we show at the end of this proof. The first integral in (A.43) involving $\partial_\alpha c_t^A - \partial_\alpha c_t^N$ is negative because $u'(c_t^A) / u'(c_T^A) > u'(c_t^N) / u'(c_T^N)$ for all $t < T$ with $T \equiv T^{\text{LF}}$, and

$$\frac{1}{2} (\partial_\alpha c_t^A - \partial_\alpha c_t^N) = \mu_t^A \partial_\alpha w_t^A - (1 - \mu_t^N) \partial_\alpha w_t^N < 0, \quad (\text{A.45})$$

as automation lowers the wages of automated workers $\partial_\alpha w_t^A = \partial_{A\alpha}^2 G^*(\cdot) < 0$ by Assumption 1, and increases the wages of non-automated workers $\partial_\alpha w_t^N = \partial_{N\alpha}^2 G^*(\cdot) > 0$ by Assumption 1 too. Moreover, the last integral in (A.43) is also negative. To see that, note that

$$\int_0^{+\infty} \omega_t^A \Delta_t^* dt < \int_0^{+\infty} \omega_t^N \Delta_t^* dt = 0. \quad (\text{A.46})$$

using Lemma 3 and the definition of ω_t^h in (A.35). The reason is that Δ_t^* increases over time and $\{\omega_t^A - \omega_t^N\}$ decreases over time. Therefore, the inequality (A.43) holds since $\tilde{\Psi}$ is sufficiently small, and thus (A.40) holds as well. This shows that the Pareto improvement requires taxing automation $\delta\alpha < 0$.⁴³

Finally, note that the inequality (4.6) in the text,

$$\int_0^T \exp(-\rho t) \left\{ \frac{u'(c_t^A)}{u'(c_T^A)} - \frac{u'(c_t^N)}{u'(c_T^N)} \right\} \partial_\alpha c_t^A dt + \tilde{\Gamma} < 0, \quad (\text{A.47})$$

follows from (A.42) up to a term

$$\tilde{\Gamma} \equiv \int_T^{+\infty} \exp(-\rho t) \left\{ \frac{u'(c_t^A)}{u'(c_T^A)} - \frac{u'(c_t^N)}{u'(c_T^N)} \right\} \partial_\alpha c_t^A dt \quad (\text{A.48})$$

that is sufficiently small by Assumption 3, as we show at the end of this proof. Put it differently, taxing automation generates a Pareto improvement whenever it increases the income (and hence consumption) of displaced workers early on $\partial\alpha \times \partial_\alpha c_t^A dt > 0$ when they are borrowing constrained and value it more, since $u'(c_t^A)/u'(c_T^A) > u'(c_t^N)/u'(c_T^N)$ for all $t < T$ with $T \equiv T^{\text{LF}}$.

Assumption 3. It remains to show that there exists a $Z^* > 0$ such that $\tilde{\Omega}$, $\tilde{\Psi}$, and $\tilde{\Gamma}$ are sufficiently small. Note that in the limit where $\theta \rightarrow 0$ and $1/\kappa \rightarrow 0$, the MRS are equalized $u'(c_t^A)/u'(c_T^A) = u'(c_t^N)/u'(c_T^N)$ once reallocation is over for $t \geq T$ with $T \equiv T^{\text{LF}}$. This immediately implies that $\tilde{\Omega} \rightarrow 0$ using the resource

⁴³ In our economy, non-automated workers and the firm agree on how they value income over time (Lemma 3). More generally, there is a Pareto improvement that requires taxing automation as soon as automated workers are effectively more impatient than non-automated workers — since the first term in (A.43) is still negative — and the firm is at least as patient as non-automated workers — since the last term in (A.43) would still be negative.

constraint (A.25) and Lemma 1, and therefore $\tilde{\Psi} \rightarrow 0$ and $\tilde{\Gamma} \rightarrow 0$ as well. By continuity of the equilibrium in θ and $1/\kappa$, there exists a $Z^* > 0$ such that $\tilde{\Omega}$, $\tilde{\Psi}$, and $\tilde{\Gamma}$ are sufficiently small for the inequalities (A.40), (A.43), and (A.47) to hold when Assumption 3 is satisfied. Finally, the threshold $Z^* > 0$ in Assumption 3 is the minimum between this one and the ones identified in Appendices A.1 and A.2.

A.5 Proof of Proposition 2

The derivative of the social welfare function with respect to automation is

$$\partial_\alpha U = \sum_h \eta^{\text{effic},h} \int_0^{+\infty} \exp(-\rho t) u'(c_t^h) \times \dot{c}_t^{h,*} dt \quad (\text{A.49})$$

when using efficiency weights $\eta^{\text{effic},h} = 1 / \int_0^{+\infty} \exp(-\rho t) u'(c_t^h) dt$, using the fact that the permanent effects are distributional $\bar{c}^{A,*} + \bar{c}^{N,*} = 0$. In this case, the intervention can *only* improve welfare through the efficiency component in (4.7).

Equivalently,

$$\partial_\alpha U = \int_0^{+\infty} (\omega_t^A - \omega_t^N) \times \partial_\alpha c_t^A dt \quad (\text{A.50})$$

$$= \int_0^{+\infty} (\omega_t^A - \omega_t^N) \times \left\{ \Delta_t^* + \mu_t^A \partial_\alpha w_t^A - (1 - \mu_t^N) \partial_\alpha w_t^N \right\} dt \quad (\text{A.51})$$

where the first equality follows from the definition of the efficiency weights, the definition of $\{\omega_t^h\}$ in (A.35), and the definition of the time-varying effects in (4.4); and the second equality follows from (A.23) and the resource constraint (A.25). We have already shown in (A.46) that

$$\int_0^{+\infty} (\omega_t^A - \omega_t^N) \times \Delta_t^* dt < 0, \quad (\text{A.52})$$

when starting from the laissez-faire. We next show that

$$\int_0^{+\infty} (\omega_t^A - \omega_t^N) \times \underbrace{\left(\mu_t^A \partial_\alpha w_t^A - (1 - \mu_t^N) \partial_\alpha w_t^N \right)}_{\equiv \Sigma_t^{A,*}} dt < 0. \quad (\text{A.53})$$

The expression (A.53) is negative as soon as $\Sigma_t^{A,*}$ increases over time. The reason is

that the sequence $\{\omega_t^A - \omega_t^N\}$ integrates to zero and decreases over time. The sequence $\Sigma_t^{A,*}$ indeed increases over time, by Assumptions 3 and 4 as we discuss at the end of this proof.

Taken together, the previous steps imply that $\partial_\alpha U < 0$ so the government finds it optimal to tax automation $\delta\alpha < 0$ locally starting from the laissez-faire α^{LF} . When the government's problem (4.1) is convex, this also implies that the laissez-faire automation is excessive compared to its second best counterpart $\alpha^{\text{SB,effic}} < \alpha^{\text{LF}}$.

Assumptions 3 and 4. It remains to show that $\Sigma_t^{A,*}$ increases over time for some $Z^* > 0$ in Assumption 3. In equilibrium $w_t^h = \partial_h G^*(\mu_t^A, \mu_t^N; \alpha)$. Therefore,

$$\begin{aligned} \Sigma_t^{A,*} = & \mu_t^A \left(\partial_{A\alpha}^2 G^*(\mu_t^A, \mu_t^N; \alpha) - \partial_{N\alpha}^2 G^*(\mu_t^A, \mu_t^N; \alpha) \right) \\ & - \left(1 - \mu_t^N - \mu_t^A \right) \partial_{N\alpha}^2 G^*(\mu_t^A, \mu_t^N; \alpha). \end{aligned} \quad (\text{A.54})$$

In the limit without productivity loss ($\theta \rightarrow 0$) or unemployment spells ($1/\kappa \rightarrow 0$), we have $\mu_t^N + \mu_t^A \rightarrow 1$. Therefore,

$$\begin{aligned} \partial_t \Sigma_t^{A,*} \rightarrow & \partial_t \mu_t^A \times \left(\partial_{A\alpha}^2 G^*(\cdot) - \partial_{N\alpha}^2 G^*(\cdot) \right) \\ & + \mu_t^A \times \partial_t \mu_t^A \left(\partial_{AA\alpha}^3 G^*(\cdot) + \partial_{NN\alpha}^3 G^*(\cdot) - 2\partial_{AN\alpha}^3 G^*(\cdot) \right) \\ = & \partial_t \mu_t^A \times \partial_{\mu\alpha}^2 \left\{ \mu \partial_\mu G^*(\mu, 1 - \mu; \alpha) \right\} \Big|_{\mu=\mu_t^A} > 0, \end{aligned} \quad (\text{A.55})$$

using Assumption 4 and the fact that $\partial_t \mu_t^A < 0$ from (3.4). By continuity of the equilibrium in θ and $1/\kappa$, there exists a $Z^* > 0$ such that $\Sigma_t^{A,*}$ increases over time. Finally, the threshold $Z^* > 0$ in Assumption 3 is the minimum between this one and the ones identified in Appendices A.1, A.2, and A.4.

A.6 Proof of Proposition 3

The law of motion of automation is $d\alpha_t = (x_t - \delta\alpha_t) dt$ where $\delta > 0$ is the depreciation rate and x_t is gross investment. Output (net of investment cost) is⁴⁴

$$Y_t = G^*(\mu_t^A, \mu_t^N; \alpha_t) - q_t x_t. \quad (\text{A.56})$$

⁴⁴ The aggregate production function $G^*(\cdot)$ does not include a cost $\mathcal{C}(\cdot)$, contrary to (5.1) in our benchmark model, since the cost of automation is now captured by $q_t x_t$.

The investment cost q_t decreases over time and converges to $q_t \rightarrow \bar{q}$ as $t \rightarrow +\infty$. Generations are indexed by s , and are born and die at rate χ . We show below that the equilibrium converges to a first best in the long-run. We refer the interested reader to the working paper version [Beraja and Zorzi \(2022\)](#) for a full description of the equilibrium with overlapping generations and the first best planning problem. We will show that the laissez-faire coincides with a particular first best allocation in the long-run.

Laissez-faire. We guess (and verify) that automation, the labor allocation and the interest rate all converge to a long-run steady state with $r_t^{\text{LF}} \rightarrow \rho$ as $t \rightarrow +\infty$. We omit the time indices at the final steady state. If the labor allocation converges to a steady state, i.e., $\mu_t^{h,\text{LF}} \rightarrow \mu^{h,\text{LF}}$ as $t \rightarrow +\infty$ in each $h = A, N$, then investment and automation also converge to steady state levels, i.e., $\alpha_t^{\text{LF}} \rightarrow \alpha^{\text{LF}}$ and $x_t^{\text{LF}} \rightarrow x^{\text{LF}}$ as $t \rightarrow +\infty$ with

$$\partial_\alpha G^* \left(\mu^{A,\text{LF}}, \mu^{N,\text{LF}}; \alpha^{\text{LF}} \right) = (\rho + \delta) \underline{q} \quad (\text{A.57})$$

and $x^{\text{LF}} = \delta \alpha^{\text{LF}}$. Conversely, if automation converges to a steady state level, so does the labor allocation and wages converge to

$$w^{A,\text{LF}} = \partial_A G^* \left(\mu^{\text{LF}}, 1 - \mu^{\text{LF}}; \alpha^{\text{LF}} \right) = \partial_N G^* \left(\mu^{\text{LF}}, 1 - \mu^{\text{LF}}; \alpha^{\text{LF}} \right) = w^{N,\text{LF}}, \quad (\text{A.58})$$

as the entry of new generations implies that the marginal products of labor (and so wages) must be equal across occupations in the long-run. Equations (A.57)-(A.58) pin down the long-run labor allocation $\{\mu^{A,\text{LF}}, \mu^{N,\text{LF}}\} = \{\mu^{\text{LF}}, 1 - \mu^{\text{LF}}\}$, automation α^{LF} , and aggregate consumption

$$C^{\text{LF}} = G^* \left(\mu^{\text{LF}}, 1 - \mu^{\text{LF}}; \alpha^{\text{LF}} \right) - \underline{q} \delta \alpha^{\text{LF}}. \quad (\text{A.59})$$

Finally, all workers are hand-to-mouth ($\underline{a} \rightarrow 0$) so $c_{s,t}^{\text{LF}} \rightarrow C^{\text{LF}}$ as $t \rightarrow +\infty$ for all generations s .⁴⁵ Therefore, $u'(c_{s,t+\tau}^{\text{LF}}) / u'(c_{s,t}^{\text{LF}}) \rightarrow 1$ as $t \rightarrow +\infty$ for all workers and horizons $\tau \geq 0$. This confirms that the interest rate $r_t^{\text{LF}} \rightarrow \rho$ as $t \rightarrow +\infty$, and the

⁴⁵ We do not index consumption by the worker's initial industry of occupation. The reason is that the mass of surviving members of old generations (born in $s < 0$) vanishes asymptotically, and new generations (born in $s \geq 0$) can choose their initial occupation of employment.

guess is verified.

First best. Proceeding as above, we can show that any first best allocation also converges to a steady state. We will show that this first best allocation is the same as the one that prevails at the laissez-faire when the planner discounts generations with the subjective discount rate ρ , i.e., it uses weights $\eta_s^h = \exp(-\rho s)$.⁴⁶ The automation choice satisfies (A.57) in the long-run when the planner discounts generations at rate ρ . Production efficiency requires that the marginal products of labor must be equalized in the long-run, so equation (A.58) holds too. Therefore, the *aggregate* allocation coincides with the laissez-faire in the long-run. It remains to show that *individual* consumptions are equal at this allocation. Note that the planner equalizes weighted marginal utilities across workers in each period t , so

$$\frac{\eta_s^h \exp(-\rho(t-s)) u'(c_{s,t}^{\text{FB}})}{\eta_{s'}^j \exp(-\rho(t-s')) u'(c_{s',t}^{\text{FB}})} = 1 \quad (\text{A.60})$$

for generations $s, s' \leq t$. Thus, consumption is equalized across workers when the planner discounts generations at rate ρ . Therefore, $c_{s,t}^{\text{FB}} \rightarrow C^{\text{FB}} = C^{\text{LF}}$ as $t \rightarrow +\infty$ for all generations s . This completes the proof.

A.7 Example

Using our example from Section 2.1, we show that an increase in the degree automation α decreases the marginal productivity of labor (MPL) *within* the automated occupation, while potentially raising the *aggregate* MPL.

The log-change in the MPL in the automated occupation is

$$\frac{d}{d\alpha} \log(\text{MPL}^A) = -\frac{1}{v} \frac{1}{y^A} \frac{(y^N)^{\frac{v-1}{v}}}{(y^A)^{\frac{v-1}{v}} + (y^N)^{\frac{v-1}{v}}} < 0.$$

Moreover,

$$\frac{d}{d\alpha} \log(\text{MPL}^N) = \frac{1}{v} \frac{1}{y^A} \frac{(y^A)^{\frac{v-1}{v}}}{(y^A)^{\frac{v-1}{v}} + (y^N)^{\frac{v-1}{v}}} > 0.$$

⁴⁶ All continuation values are evaluated at birth as in Calvo and Obstfeld (1988).

That is, the MPL declines in the automated occupation but increases in the non-automated occupation. The marginal productivity of labor at the aggregate level, i.e., workers' average wage rate, is

$$\text{MPL} \equiv \frac{\mu^A}{\mu^A + \mu^N} \text{MPL}^A + \frac{\mu^N}{\mu^A + \mu^N} \text{MPL}^N,$$

and it can increase or decrease, depending on (μ^A, μ^N, ν) .

A.8 An Economy Where Automated Workers Claim No Profits

Consider an economy where automated and non-automated workers trade equity, but non-automated workers hold all the equity initially. We rule out borrowing and short sales $\underline{a} \rightarrow 0$, as in Section 4. Therefore, automated workers A do not claim *any* profits. We next show that a Pareto improvement still exists in this economy, and it requires taxing automation (as in Proposition 1). Taxing automation is also still optimal on efficiency grounds (as in Proposition 2).

To show that there is a Pareto improvement and that it requires taxing automation (Proposition 1), we follow the exact same steps as in Appendix A.4. All steps (A.37)–(A.41) remain the same, and $\tilde{\Omega}$ is still sufficiently small. Therefore, inequality (A.42) holds or, proceeding as in Appendix A.4,

$$\int_0^T \exp(-\rho t) \left\{ \frac{u'(c_t^A)}{u'(c_T^A)} - \frac{u'(c_t^N)}{u'(c_T^N)} \right\} \partial_\alpha c_t^A dt < 0 \quad (\text{A.61})$$

up to a term that is sufficiently small. The intuition is unchanged: taxing automation generates a Pareto improvement whenever it increases the income (and hence consumption) of displaced workers early on $\partial_\alpha \times \partial_\alpha c_t^A dt > 0$ when they are borrowing constrained and value it more, since $u'(c_t^A) / u'(c_T^A) > u'(c_t^N) / u'(c_T^N)$ for all $t < T$.

The inequality (A.61) holds as soon as automation makes automated workers worse off while they reallocate (before period T), i.e., $\partial_\alpha c_t^A < 0$ for $t < T$. The only difference compared to the benchmark model is that the change in income (and hence consumption) of automated workers after a perturbation in automation $\delta\alpha$ now works exclusively through labor income (since they do not claim profits)

$$\partial_\alpha c_t^A = 2 \times \left(\mu_t^A \partial_\alpha w_t^A + \left(\mu_t^N - \frac{1}{2} \right) \partial_\alpha w_t^N \right). \quad (\text{A.62})$$

In particular, taxing automation $\delta\alpha < 0$ lifts the wage of automated workers early on $\delta\alpha \times \partial_\alpha w_t^A > 0$ when automated workers value consumption the most.

The only scenario where automation makes automated workers better off early on, i.e., $\partial_\alpha c_t^A$ turns positive before T , is a rather pathological one. It must be that automation raises the wage in the non-automated occupation w_t^N *so much* that the average income of displaced workers increases — as workers reallocate over time to the N occupation.⁴⁷ This is a rather uninteresting case in our context, and it is also counterfactual (Acemoglu and Restrepo, 2022; Braxton and Taska, 2023).

Finally, we show that taxing automation remains optimal on efficiency grounds alone (Proposition 2). The derivative of the social welfare function with respect to automation is the same as (A.50). Therefore, the government finds it optimal to tax automation on efficiency grounds when

$$\partial_\alpha U = \int_0^{+\infty} \left(\omega_t^A - \omega_t^N \right) \times \partial_\alpha c_t^A dt < 0. \quad (\text{A.63})$$

This derivative is again negative as soon as automation depresses the income (and hence consumption) of automated workers more early on (before they have time to reallocate), i.e., $\partial_\alpha c_t^A$ increases over time. This is the case under Assumptions 3 and 4 when $\partial_{\alpha\mu}^2 \{ \partial_N G^*(\mu, 1 - \mu; \alpha) \}$ is either negative or sufficiently small.⁴⁸ In this case, taxing automation is optimal when using the efficiency weights $\eta^{h,\text{effic}}$.

A.9 An Alternative Decomposition

In expression (4.7), we decompose changes in welfare between those associated with time-varying or permanent effects of automation on consumption. We label the former “efficiency” and the latter “equity.” The reason is the following. In an efficient economy where the intertemporal MRS of all workers coincide with the equilibrium interest rate, a perturbation in automation does not affect welfare

⁴⁷ This pathological case does not arise in our benchmark model. The reason is that an increase in the wage w_t^N also reduces profits claimed by all workers.

⁴⁸ For instance, $\partial_{\alpha\mu}^2 \{ \partial_N G^*(\mu, 1 - \mu; \alpha) \} = 0$ in the example in Section 2.1 when evaluated in a symmetric allocation with $y^A = y^N$.

through the time-varying effects. In this case, the efficiency term is zero. All welfare changes come from permanent effects, which are purely distributional. Our decomposition captures this precisely through the equity component. In an inefficient economy where borrowing constraints bind, a perturbation in automation now also affects welfare through the time-varying effects. Our efficiency component precisely captures all the consumption changes that can improve welfare *only* when the economy is inefficient to begin with. In particular, our decomposition associates Pareto improvements with efficiency improvements, since Pareto improvements are only possible when taxing automation has time-varying effects — see equation (A.39) and the related discussion.

An alternative approach is to decompose welfare changes based on whether they come from aggregate changes in consumption or distributional changes across workers, regardless of whether they improve efficiency or equity (as we have defined them). That is,

$$\begin{aligned}\text{Aggregate} &\equiv \sum_h \eta^h \int_0^{+\infty} \exp(-\rho t) u'(c_t^h) \times \Delta_t^* dt \\ \text{Redistribution} &\equiv \sum_h \eta^h \int_0^{+\infty} \exp(-\rho t) u'(c_t^h) \times \left\{ \partial_\alpha c_t^h - \Delta_t^* \right\} dt\end{aligned}$$

in the case of automation. This decomposition is similar to the one in [Bhandari et al. \(2021\)](#). Both our decomposition (4.7) and this alternative one are valid and mutually consistent. They just capture different ideas.

To see the connection between these two decompositions, consider the welfare change from taxing automation when using efficiency weights $\eta^{h,\text{effic}}$ starting from the laissez-faire. As we have shown in Appendix A.5, taxing automation improves welfare *only* through the efficiency component in (4.7). At the same time, equations (A.52) and (A.53) show that this increase in efficiency can be further decomposed into two effects. The first effect in (A.52) is exactly the same as the “aggregate” effect in the alternative decomposition. The second effect in (A.53) is instead captured by the “redistribution” effect in the alternative decomposition. With this in mind, one can say that taxing automation improves efficiency both because it increases aggregate output and redistributes income early on during the transition, precisely when displaced workers become borrowing constrained. Both of these effects would be zero in an efficient economy. In our economy, they create a motive

to intervene on efficiency grounds alone, which is distinct from the equity motive put forth in the literature (Guerreiro et al., 2022; Costinot and Werning, 2022).

B Quantitative Model

In this appendix, we describe our quantitative model in more detail. Section B.1 provides a recursive formulation of the workers' problem. Section B.2 states and characterizes the solution to the firm's problem. Section B.3 discusses the second best.

B.1 Workers' Problem

We discretize time into periods of constant length $\Delta \equiv 1/N > 0$, and solve the workers' problem in discrete time.⁴⁹ The workers' problem can be formulated recursively:

$$\begin{aligned} V_t^h(a, e, \xi, z) &= \max_{c, a'} u(c) \Delta + \exp(-(\rho + \chi) \Delta) V_{t+\Delta}^{h,*}(a', e, \xi, z) \\ \text{s.t. } a' &= (\mathcal{Y}_t^{\text{net}}(\mathbf{x}) - c) \Delta + \frac{1}{1 - \chi \Delta} (1 + r_t \Delta) a \\ a' &\geq 0 \end{aligned} \quad (\text{B.1})$$

for employed workers ($e = E$) and unemployed workers ($e = U$). The continuation value V^* before workers observe the mean-reverting component of their income is given by

$$V_t^{h,*}(a, e, \xi, z) = \int \hat{V}_t^h(a, e, \xi, z') P(dz', z), \quad (\text{B.2})$$

where $\hat{V}_t(\cdot)$ is the continuation value associated to the discrete occupational choice. The continuation value for employed workers ($e = E$) associated to this discrete

⁴⁹ Alternatively, we could have formulated the workers' problem in continuous time and solved the associated partial differential equation using standard finite difference methods. However, (semi-)implicit schemes are non-linear in our setting due to the discrete occupational choice. This requires iterating on (B.1)–(B.5) to compute policy functions which limits the efficiency of these schemes. We found that explicit schemes were unstable unless we use a particularly small time step Δ which is again relatively inefficient. Formulating and solving the workers' problem in discrete time proves to be relatively fast.

choice problem is⁵⁰

$$\hat{V}_t^h(a, e, \xi, z) = (1 - \lambda \Delta) V_t^h(a, e, \xi, z) + \lambda \Delta \gamma \log \left(\sum_{h'} \phi^{h'} \exp \left(\frac{V_t^{h'}(a, e'(h', \mathbf{x}), \xi, z)}{\gamma} \right) \right), \quad (\text{B.3})$$

with $e'(\cdot) = E$ if $h' = h$ and $e'(\cdot) = U$ otherwise. The associated mobility hazard across occupations is

$$\mathcal{S}_t(h'; \mathbf{x}) = \frac{\phi^{h'} \exp \left(\frac{V_t^{h'}(\mathbf{x}'(h'; \mathbf{x}))}{\gamma} \right)}{\sum_{h''} \phi^{h''} \exp \left(\frac{V_t^{h''}(\mathbf{x}'(h''; \mathbf{x}))}{\gamma} \right)}, \quad (\text{B.4})$$

where $\mathbf{x}'(h'; \mathbf{x})$ is short for $(a, e'(h', \mathbf{x}), \xi, z)$. In turn, the continuation value for unemployed workers ($e = U$) is

$$\hat{V}_t^h(a, e, \xi, z) = (1 - \kappa \Delta) V_t^h(a, e, \xi, z) + \kappa \Delta V_t^h(a, 1, \xi'(h', \mathbf{x}), z), \quad (\text{B.5})$$

where $\xi'(\cdot) = (1 - \theta) \xi$ after the first reallocation spell is complete. New generations who enter the labor market draw a random productivity z from its stationary distribution and then choose their occupation with a hazard similar to the employed workers'. The only difference is that they experience neither an unemployment spell nor a productivity loss. Worker's gross labor income (or unemployment benefits) is

$$\mathcal{Y}_t^{\text{labor}}(\mathbf{x}) = \begin{cases} \xi \exp(z) w_t^h & \text{if } e = E \\ b \mathcal{Y}_t^{h'}(a, E, \xi, z) & \text{otherwise} \end{cases}, \quad (\text{B.6})$$

with $h' \neq h$ denoting the previous occupation of employment. The permanent component of workers' income (ξ) is reduced by a factor $(1 - \theta)$ whenever a worker who exits unemployment enters her new occupation. Workers experience this productivity loss at most once during their lifetime. Finally, the mean-reverting component of income (z) evolves as

$$z' = (1 + (\tilde{\rho}_z - 1) \Delta) z + \sigma_z \sqrt{\Delta} W' \quad \text{with} \quad W' \sim \text{i.i.d.} \mathcal{N}(0, 1), \quad (\text{B.7})$$

⁵⁰ See Artuç et al. (2010) for the derivation.

where $\tilde{\rho}_z \equiv 1 - \rho_z$ governs its persistence.

B.2 Firm's Problem

We solve the firm's problem in continuous time. The problem is

$$\begin{aligned} \max_{\{x_t, \alpha_t, \mu_t^A, \mu_t^N\}} \int_0^{+\infty} \exp\left(-\int_0^t r_s ds\right) & \left\{ G^*(\mu_t^A, \mu_t^N; \alpha_t) - \phi q_t x_t - \sum_h \phi^h w_t^h \mu_t^h \right\} dt \\ \text{s.t. } d\alpha_t &= (x_t - \delta \alpha_t) dt, \alpha_0 = 0, x_t \geq 0 \end{aligned} \quad (\text{B.8})$$

where α_t is the stock of automation, x_t is gross investment, μ_t^A and μ_t^N are labor demands, q_t is the resource cost per unit of investment, and $G^*(\mu_t^A, \mu_t^N; \alpha_t) \equiv G(y_t^A, y_t^N)$ with y_t^A and y_t^N given by (5.1).⁵¹ The optimal degree of automation satisfies

$$(r_t + \delta) \phi q_t = \partial_\alpha G^*(\mu_t^A, \mu_t^N; \alpha_t) + \phi \partial_t q_t, \quad (\text{B.9})$$

together with the law of motion

$$d\alpha_t = (x_t - \delta \alpha_t) dt, \quad (\text{B.10})$$

and the initial condition $\alpha_0 = 0$. Finally, the firm's labor demands satisfy

$$w_t^h = (1 - \eta) \frac{1}{\alpha_t^h + \mu_t^h} \frac{\left\{ A^h (\alpha_t^h + \mu_t^h)^{(1-\eta)} \right\}^{\frac{\nu-1}{\nu}}}{\sum_g \phi^g \left\{ A^g (\alpha_t^g + \mu_t^g)^{(1-\eta)} \right\}^{\frac{\nu-1}{\nu}}} G^*(\mu_t^A, \mu_t^N; \alpha), \quad (\text{B.11})$$

where $\alpha_t^A \equiv \alpha_t$ and $\phi^A \equiv \phi$ in automated occupations and $\alpha_t^N \equiv 0$ and $\phi^N \equiv 1 - \phi$ in non-automated occupations, and

$$\mu_t^h = \frac{1}{\phi^h} \int \mathbf{1}_{\{e(\mathbf{x})=1, h(\mathbf{x})=h\}} \xi \pi_t(d\mathbf{x}) \quad (\text{B.12})$$

⁵¹ For concision, we omit any distortionary tax on investment τ_t^x since it is isomorphic to the cost of investment q_t . The aggregate production function $G^*(\cdot)$ does not include a cost $\mathcal{C}(\cdot)$ as discussed in footnote 44.

is the (effective) labor supplied in *each* occupation.

B.3 Second Best

In this appendix, we state the second best problem that we solve numerically and discuss our choice of Pareto weights.

Objective. The government's objective is

$$\begin{aligned} \mathcal{W} \equiv & \chi \int_{-\infty}^0 \int \eta_s(\mathbf{x}) \exp(-(\rho + \chi)(0 - s)) V_0^{\text{old}}(\mathbf{x}) \pi_{s,0}^{\text{old}}(d\mathbf{x}) ds \\ & + \chi \int_0^{+\infty} \eta_s V_s^{\text{new}} ds, \end{aligned} \quad (\text{B.13})$$

for some Pareto weights η , where $\pi_{s,0}^{\text{old}}$ is the initial distribution of idiosyncratic states for existing generations born in $s < 0$ (conditional on survival). Following [Calvo and Obstfeld \(1988\)](#), all continuation values are evaluated at birth. The value $\exp(-(\rho + \chi)(0 - s)) V_0^{\text{old}}$ is the continuation utility of existing generations over periods $t \geq 0$. In turn, the value

$$V_t^{\text{new}} \equiv \int \gamma \log \left(\sum_h \phi^h \exp \left(\frac{V_t^h(0, 1, 1, z)}{\gamma} \right) \right) P^*(dz) \quad (\text{B.14})$$

is the continuation utility for new generations born in period $t = s \geq 0$, which reflects their occupational choice.⁵² Here, P^* denotes the ergodic distribution of the income process $z'|z \sim P(z)$, i.e., the distribution of productivities at birth.

Pareto weights. We use *efficiency* weights that capture the efficiency motive for policy intervention and ensure that the government has no incentive to redistribute resources. These weights are the ones we described in our baseline model (Section 4.5). The government discounts generations using the subjective discount rate ρ as

⁵² Members of a new generation are born with no assets $a = 0$, are employed $e = 1$, and have not incurred the productivity cost associated to switching occupations $\xi = 1$.

in [Itskhoki and Moll \(2019\)](#) and [Guerreiro et al. \(2022\)](#). Therefore, the weights are

$$\eta_s(\mathbf{x}) = \frac{\exp(-\rho s)}{\int_0^{+\infty} \exp(-(\rho + \chi)t) u'(C_t^{\text{old},h}) dt} \quad (\text{B.15})$$

for old generations employed in automated ($h = A$) and non-automated ($h = N$) occupations, where $C_t^{\text{old},h}$ denotes average consumption over time for each of these groups at the laissez-faire.⁵³ The weights for new generations (indexed by $s > 0$) are similar, except that they are not indexed by h since new generations are able to choose their initial occupation of employment. They depend on the consumption streams $C_{s,t}^{\text{new}}$ for all $t \geq s$.

Summarizing, the government's objective becomes

$$\begin{aligned} \mathcal{W} \equiv & \int \frac{V_0^{\text{old}}(\mathbf{x})}{\int_0^{+\infty} \exp(-(\rho + \chi)t) u'(C_t^{\text{old},h}) dt} \pi_0(d\mathbf{x}) ds \\ & + \chi \int_0^{+\infty} \exp(-\rho s) \frac{V_s^{\text{new}}}{\int_s^{+\infty} \exp(-(\rho + \chi)(t-s)) u'(C_{s,t}^{\text{new}}) dt} ds, \end{aligned} \quad (\text{B.16})$$

where

$$\pi_0(d\mathbf{x}) \equiv \int_{-\infty}^0 \chi \exp(\chi s) \pi_{s,0}^{\text{old}}(d\mathbf{x}) ds \quad (\text{B.17})$$

is the initial distribution of idiosyncratic states.

Policy tools and implementability. The government maximizes the objective (B.16) by choosing an appropriate sequence of distortionary taxes on investment $\{\tau_t^x\}$ and rebating the proceedings to the firm. The implementability constraints consist of workers' reallocation and consumption choices, the labor market clearing conditions, and the aggregate resource constraint.

⁵³ An alternative approach would be to compute the expected marginal utility over time separately for each initial state $(\mathbf{x})$. This is computationally infeasible since our state space is too large.

C Numerical Implementation

We discuss how we solve numerically for the equilibrium and the optimal policy.

Workers' problem. We solve the worker's problem (B.1)–(B.5) using the standard endogenous grid method (Carroll, 2006). In theory, this problem could be non-convex since it involves a discrete choice across occupations. However, we find that this is not the case in our calibration. The variance of the taste shocks γ is sufficiently large that the value function remains concave. We use Young (2010)'s non-stochastic simulation method to iterate on the distribution. Finally, we discretize the income process on a 7-point grid using the method of Rouwenhorst (1995).

Firm's problem. Given a sequence for the investment cost $\{q_t\}$ and the interest rate $\{r_t\}$, the optimal sequence of automation and investment can be solved using (B.9)–(B.10) with initial condition $\alpha_0 = 0$. The initial cost of investment ensures that automation is continuous in $t = 0$ at the laissez-faire

$$q_0 = \frac{1}{\phi} \frac{\partial_\alpha G^*(\mu_0^A, \mu_0^N; \alpha_0) + \phi q_\infty}{r_0 + \delta + \phi}, \quad (\text{C.1})$$

using $q_t = q_\infty + \exp(-\phi t)(q_0 - q_\infty)$ where $q_\infty \equiv \lim_{t \rightarrow +\infty} q_t$ is the long-run cost of investment.

Policy. We adopt a primal approach as in our benchmark model. We assume that the government can directly choose the wage gap across occupations $\hat{w}_t \equiv \log(w_t^A) - \log(w_t^N)$. For numerical reasons, we restrict our attention to parametric perturbations. Specifically, we consider policies of the form

$$\hat{w}_t = S(t; \Theta) \hat{w}_t^{\text{LF}} \quad (\text{C.2})$$

where $\hat{w}_t^{\text{LF}}$ is the wage gap that prevails at the laissez-faire, and

$$S(t; \Theta) \equiv \min \left\{ \max \left\{ 3 \left(\frac{t}{\Theta} \right)^2 - 2 \left(\frac{t}{\Theta} \right)^3, 0 \right\}, 1 \right\} \quad (\text{C.3})$$

is a smoothstep function with argument t and scale parameter Θ .⁵⁴ We search for the optimal Θ over a grid, computing welfare (B.16) for each point. The second best intervention is the one that delivers the highest welfare.

We proceed as follows to recover the taxes on investment $\{\tau_t^x\}$ that implement the second best. We define $q_t^{\text{SB}} \equiv (1 + \tau_t^x) q_t$, where q_t is the laissez-faire cost. Using (B.9),

$$(r_t + \delta) \phi q_t^{\text{SB}} = \partial_\alpha G^* \left(\mu_t^{A,\text{SB}}, \mu_t^{N,\text{SB}}; \alpha_t^{\text{SB}} \right) + \phi \partial_t q_t^{\text{SB}}, \quad (\text{C.4})$$

where the allocations are evaluated at the second best. This expression defines a differential equation for $\{q_t^{\text{SB}}\}$ with terminal condition $\lim_{t \rightarrow +\infty} q_t^{\text{SB}}/q_t = 1$ since the second best allocation converges to its laissez-faire level. We solve this differential equation using a standard shooting algorithm, and we recover the taxes $\tau_t^x = q_t^{\text{SB}}/q_t - 1$.

Welfare gains. We compute the welfare gains as the ratio between the certainty equivalent consumptions that produce the same welfare as in the second best and the laissez-faire, respectively. The welfare gains are thus given by $(\mathcal{W}^{\text{SB}}/\mathcal{W}^{\text{LF}})^{\frac{1}{1-\sigma}}$, where $\sigma > 0$ is the inverse elasticity of intertemporal substitution.

D Output Share and Employment

We argued that our model matches well the share of output produced by automation forecasted by McKinsey (2017), as well as the firm-level effects of automation on employment estimated by Bonfiglioli et al. (2022). We now explain how we compute the model analogs of these (untargeted) moments.

Output share. Exhibit E3 in McKinsey (2017) finds that roughly 70% of output previously produced by labor could be automated in occupations most susceptible to automation (making up for 51% of initial employment, compared to 56% in our model). This figure is obtained by taking the weighted average of the

⁵⁴ Note that these policies constrain the government to intervene only along the transition. In theory, the government might also want to intervene in the long-run due to uninsured idiosyncratic risk (Dávila et al., 2012). By construction, allowing for such, more flexible, policy would produce even higher welfare gains compared to Table 6.2. We chose to abstract from long-run taxation to focus on the new motive for intervention that we highlight (Section 4.6.2).

time spent on automatable activities in the three most susceptible activities $0.71 = (17 \times 64 + 16 \times 69 + 18 \times 81) / (17 + 16 + 18)$. In our model, the share of output in occupation $h = A$ that is produced by automation is $\alpha / (\alpha + \mu^A)$, which is 67% when evaluated at the final steady state.

Employment. The percent change in employment of a firm that adopted automation, relative to a firm that did not, can be computed by the ratio of the coefficients in column (2) to column (5) in the first line of Table 2 of Bonfiglioli et al. (2022). This gives $-0.094/0.174 = -54\%$. In our model, labor demand from a “firm” producing the intermediate good supplied by automated occupations satisfies

$$A (1 - \eta) (\alpha + \mu)^{-\eta} = \frac{w}{p}, \quad (\text{D.1})$$

where w is wage and p is the price of the intermediate good produced in automated occupations. Next, we consider the following partial equilibrium exercise. Let us compare two firms producing the same intermediate good and facing the same wage and price. One is partially automated $\alpha_1 > 0$ and the other one is not $\alpha_0 = 0$. Then, it must be that


$$\alpha_1 + \mu_1 = \mu_0. \quad (\text{D.2})$$

So, the log-change in employment is

$$\log(\mu_1) - \log(\mu_0) = \log\left(1 - \frac{\alpha_1}{\mu_0}\right) = -0.66 \quad (\text{D.3})$$

using our calibration, where μ_0 is the initial steady state employment in automated occupations and α_1 equals the stock of automation 20 years out in the transition, which is the sample period in Table 2 of Bonfiglioli et al. (2022). Next, we consider employment changes in general equilibrium across steady states. We find that employment in automated occupations changes by -0.37 log points. Overall, our partial and general equilibrium exercises deliver predictions that are comparable to the -0.54 log change that Bonfiglioli et al. (2022) estimate.

Figure E.1: Relative income, consumption, and assets for A workers



Notes: This figure plots the average income, consumption, and assets for workers initially employed in automated occupations, in percentage deviations from those of workers initially employed in non-automated occupations.


E Income, Consumption, and Savings Dynamics

Figure E.1 plots the average income, consumption, and assets of workers initially employed in automated occupations, in percentage deviations from those of workers initially employed in non-automated occupations. The income of automated workers falls during the very first periods, as they have not had time to reallocate yet. Over time, their income profile becomes upward sloping as they are able to reallocate. This partial recovery takes place over roughly three decades. Automated workers borrow to smooth consumption while their income is low.

F Taxes on Automation

Figure F.1 plots the sequence of taxes on investment $\{\tau_t^x\}$ that implement the second best allocation, for each of the four calibrations in Table 6.2. We have discussed how we solve for these taxes in Appendix C. By construction, taxes converge to zero in the long-run (footnote 54).

Figure F.1: Investment taxes at the second best



Notes: The four curves correspond to the calibrations in Table 6.2.