Occupation Dynamics in General Equilibrium*

Benny Kleinman†
Stanford University and NBER
Ernest Liu‡
Princeton University and NBER
Stephen J. Redding§
Princeton University, NBER and CEPR
Xiang Zhang¶
Princeton University

March, 2025

Abstract

We develop a tractable quantitative model of automation and non-neutral technical change with dynamic labor mobility and capital accumulation. Given data on an economy somewhere along the transition path to an unobserved steady-state, we show how to solve for this unobserved steady-state and the counterfactual transition path in response to shocks. We invert the equilibrium conditions of the model to recover shocks that exactly rationalize the observed employment and wage data. We find relatively rapid convergence to steady-state, such that initial conditions are relatively unimportant in explaining the observed employment and wage dynamics. We find that intermediate-wage occupations have experienced the greatest automation, whereas high-wage occupations have witnessed the greatest productivity growth. We find that the combination of automation and productivity growth explains most of the observed employment and wage polarization in the data. We compute an inequality-output frontier, which summarizes the impact of automation and productivity growth on wage inequality and aggregate output, and can be used to assess the effects of emergent new technologies, such as artificial intelligence.

Keywords: automation, artificial intelligence, employment polarization, wage inequality

JEL Classification: F14, F15, F50

*We are grateful to Princeton and Stanford Universities for research support. We would like to thank Chenming Zhang for excellent research assistance. The usual disclaimer applies.

†Dept. Economics, 579 Jane Stanford Way, Stanford, CA 94305. Email: bennyk@stanford.edu.

‡Dept. Economics, JRR Building, Princeton, NJ 08544. Email: ernestliu@princeton.edu.

§Dept. Economics and SPIA, JRR Building, Princeton, NJ 08544. Email: reddings@princeton.edu.

¶Dept. Economics, JRR Building, Princeton, NJ 08544. Email: xiangzhang@princeton.edu.

1 Introduction

A key research and public policy question is the effects of new technologies, such as automation and artificial intelligence, on income distribution and welfare. These effects are likely to only emerge gradually over time, because it takes time for workers to reallocate in response to labor demand shocks, and it takes time for capital to be accumulated to realize the potential of these new technologies. These two sources of dynamics interact with one another to shape the impact of new technologies on income distribution and welfare. On the one hand, if workers can only reallocate gradually, the immediate impact of a negative shock to labor demand in an occupation can be larger, and the welfare gains from reallocation to other occupations to take advantage of the potential of the new technology can be smaller. On the other hand, if firms can only gradually accumulate capital that is complementary to new technologies, this dampens their impact on labor demand, thereby providing workers more time to adjust, and mitigating the impact of these new technologies on income distribution and welfare.

Analyzing two-sided dynamics in empirically-realistic settings with many occupations can be challenging. The speed at which workers choose to reallocate between occupations depends on expected investments in these occupations, while these investments themselves depend on the expected speed of worker reallocation. Furthermore, decisions to work and invest in any one occupation depend on expected decisions to work and invest in all other occupations. With a high-dimensional state space of many occupations, this two-sided dynamic optimization problem can become intractable. For this reason, existing research has often limited the dimensions of the state space, by either focusing on a small number of worker groups (e.g., skilled versus unskilled workers), allowing dynamics only on the worker side (e.g., dynamic discrete choice labor supply models), or considering discrete technology adoption decisions (e.g, dynamic discrete choices of whether to adopt a new technology).

We develop a quantitative theoretical model with dynamic labor mobility and capital accumulation that remains tractable even with a high-dimensional state space. At each point in time, workers supply labor in their existing occupation. They also make forward-looking decisions whether to incur mobility costs to reallocate to other occupations, subject to idiosyncratic shocks to these mobility costs. Firms produce output by performing a set of procedures that are associated with occupations, where each procedure involves tasks that can be performed either by labor or capital. Capital and labor are perfect substitutes at the task level, after adjusting for idiosyncratic shocks to the productivity of capital and labor in performing individual tasks. Aggregating across tasks, we obtain a production technology for each procedure that takes the constant elasticity of substitution (CES) form. Capital owners make forward-looking decisions about investment in each procedure subject to adjustment costs.

The state variables of the model are the share of workers and the stock of capital in each occupation. The dynamics of the state variables depends on the structural parameters of the model and fundamental shocks. The model has five key structural parameters. The degree of dispersion of idiosyncratic mobility shapes the elasticity of labor supply to each location. The elasticity of substitution between capital and labor at the procedure level governs the response of factor demand to automation changes and changes in relative factor prices. The elasticity of substitution across procedures controls reallocation across procedures in response to non-neutral technical change and changes in relative prices. The size of adjustment costs regulates the speed with the stock of capital responds to changes in investment profitability. The discount rate influences the dynamics of labor supply (through mobility across occupations) and labor demand (through capital accumulation).

The model features five key fundamental shocks that vary across procedures and over time. The automation of tasks within procedures shapes the capital-labor ratio for each procedure. Non-neutral technical change determines that reallocation of resources across occupations with different initial capital-labor ratios and wage levels. Non-wage compensating differentials capture variation in the relative attractiveness of occupations that cannot be explained by wages or continuation values. Bilateral mobility costs summarize factors that are common to all workers and influence the ease with which workers can move between occupations. Investment efficiency controls the conversion of consumption goods into capital goods and influences the relationship between the flow of investment and the accumulation of the capital stock.

Given observed data on wages, employment, capital payments, capital stocks and worker mobility for an economy that is somewhere along the transition path to a steady-state, we invert the equilibrium conditions of the model and recover unique values of the fundamental shocks that exactly rationalize the observed data as an equilibrium of the model. We do not impose any restrictions on the future evolution of fundamentals except for the assumption of perfect foresight. We show that dynamic exact-hat algebra techniques can be used to solve for the counterfactual transition path of the economy in the absence of any fundamental shocks (convergence to an initial steady-state) or in response to any convergent sequence of future fundamental shocks. If we undertake a counterfactual allowing for convergence to steady-state and all fundamental shocks, the counterfactual predictions of the model exactly reproduce the observed data in all subsequent time periods.

We use these properties of our model to measure the fundamental shocks and quantify their contributions towards the polarization of employment and wages observed in the United States over our sample period from 1981-2018. We undertake counterfactuals allowing only convergence to the initial steady-state and then successively adding fundamental shocks. Consistent with existing reduced-form empirical evidence, we find that intermediate-wage occupations have

experienced the greatest automation. In line with the literature on skill-biased technological change, we find that high-wage occupations have experienced the greatest non-neutral technical change. Although existing empirical research has sometimes considered these phenomena separately, we find that the combination of the two is required account quantitatively for observed employment and wage polarization. In comparison, non-wage compensating differentials, investment efficiency and mobility costs make relatively minor contributions to observed changes in employment and wages across occupations.

We use our model counterfactuals to assess the extent to which gains in aggregate output from technological change come at the expense of increases in wage inequality. In general, automation in low-wage occupations (bottom 20 percent of the initial wage distribution) increases wage inequality, whereas automation in high-wage occupations (top 20 percent of the initial wage distribution) reduces wage inequality. However, these is considerable heterogeneity within each of these groups in the increase in aggregate output achieved by a given amount of automation, in part related to the initial employment size of the occupations. Automation in some occupations brings large gains in aggregate output with little increase in wage inequality, while automation in other occupations has only a modest impact on aggregate output at the expense of a substantial increase in wage inequality. We find similar heterogeneity across occupations in the extent to which there is trade-off between higher aggregate output and greater wage inequality in response to non-neutral technological change. To the extent that policy makers are concerned about equity, these empirical findings provide guidance as to how policies intended to influence the direction of automation or technical change can achieve equity objectives at relatively low cost in terms of aggregate output.

To characterize the determinants of the economy's dynamics, we linearize the model around this unobserved steady-state to derive a first-order closed-form solution for the economy's transition path. We show that the economy's dynamics are fully characterized to first-order by an impact matrix that captures the initial impact of shocks on the state variables and a transition matrix that governs the updating of the state variables over time. We use this linearization to decompose the evolution of employment and wages in each occupation into the contribution of convergence given initial conditions and automation shocks. We show that the evolution of the state variables and the speed of convergence (as measured by the half-life of convergence to steady-state) can be characterized solely using the eigenvalues and eigenvectors of the transition matrix. We find a relatively rapid speed of convergence to steady-state, which is driven by the substantial worker mobility between occupations observed in our data. But we find that the speed of convergence exhibits substantial heterogeneity depending on the incidence of shocks across occupations, with half lives of convergence to steady-state ranging from a few years to decades. These findings provide guidances as to which types of automation and technology shocks the

economy can easily adjust to and which involve prolonged periods of adjustment.

Our paper is related to several existing strands of research. First, our paper contributes to the literature on directed technological change and the race between technology and skill, including Acemoglu (2002), Autor et al. (2008), Goldin and Katz (2010), Hanlon (2015), Acemoglu et al. (2015), Hémous and Olsen (2021), and Loebbing (2022), as reviewed in Acecmoglu and Autor (2011). Much of the theoretical literature on directed technological change has focused on the distinction between skilled and unskilled workers. We allow for non-neutral technological change and the accumulation of complementary physical capital across a large number of occupations. We incorporate transition dynamics from both worker mobility and capital accumulation. We develop a tractable framework that lends itself to quantitative analysis in empirical settings with large numbers of occupations.

Second, our research is related to studies of the labor market implications of new technologies, such as automation, robots and artificial intelligence, including Autor et al. (2003), Autor and Dorn (2013), Michaels and Graetz (2018), Acemoglu and Restrepo (2019), Acemoglu and Restrepo (2020), Humlum (2021), Acemoglu and Restrepo (2022) and Costinot and Werning (2023). While some of these papers focus on discrete technology adoption decisions, we model the continuous decision of the accumulation of complementary physical capital in response to automation shocks. Our quantitative model allows us to assess the relative importance of different fundamental shocks (e.g., automation, non-neutral technical change, non-wage compensating differentials) in the data. Through modelling the dynamics of labor supply and capital accumulation, we analyze the speed with which the economy adjusts to automation and non-neutral technical change.

Second, our work contributes to research on dynamic discrete choice models of labor mobility, including Artuç et al. (2010), Dix-Carneiro (2014), Caliendo et al. (2019), Traiberman (2019) and Kleinman et al. (2023). We incorporate dynamics for both worker mobility and capital accumulation. While Kleinman et al. (2023) consider a single final good produced using labor and physical capital, we models workers employed in many occupations, and distinguish between automation and non-neutral technical change. We use this distinction to explore the speed with which the economy adjusts to automation and non-neutral technical change over time.

Fourth, our analysis relates to recent debates about the labor market consequences of artificial intelligence, including Acemoglu et al. (2022), Babina et al. (2022), Acemoglu and Johnson (2023), and McElheran et al. (2023). Given observed data on wages, employment, labor-shares, and worker transitions between occupations, we show how our framework can be used to make quantitative predictions about the future evolution of wages and employment in response to new technologies, such as the emergence of artificial intelligence.

The remainder of the paper is structured as follows. Section 2 develops our theoretical framework. Section 3 develops the technical tools to study the dynamics of worker mobility and firm

investment in complementary physical capital. Section 4 describes the data and presents our main empirical results. Section 5 concludes.

2 Theoretical Framework

We consider an economy with a single final good that is produced using a number of procedures. Each of these procedures is associated with an occupation $o \in \{1, \ldots, \mathbb{O}\}$ and involves tasks that either can be performed by labor or can be automated and performed by capital. Time is discrete and is indexed by t. We use bold math font to denote vectors or matrices.

There are two types of infinitely-lived agents: workers and capitalists. Workers are endowed with one unit of labor that is supplied inelastically and can move between occupations subject to mobility costs. Workers do not have access to an investment technology, and hence live "hand to mouth," as in Kaplan and Violante (2014). Workers make dynamic discrete choice decisions between occupations based on expected continuation values and bilateral mobility frictions.

A representative capitalist owns capital and makes forward-looking decisions about investment to perform automated tasks in each occupation. There is a single integrated capital market. But investments in each occupation are subject to adjustment costs. Therefore, the rental rate can differ across occupations along the transition path to steady-state, and automation and other shocks lead to gradual capital accumulation or decumulation.

The main new elements in the model are the combination of discrete choice worker mobility decisions across many occupations (dynamic labor supply) and continuous investment in complementary capital for automated tasks (dynamic labor demand). We show how these two sources of dynamics interact to shape the impact of automation and non-neutral technical change within a single integrated framework. In our baseline specification, agents have perfect foresight and there is no aggregate uncertainty. While for simplicity, we develop the model with a single final good, our framework can accommodate multiple final goods sectors. Online Appendix B contains the derivation of all results in this section.

2.1 Worker's Problem

Utility for a worker in occupation o at time t depends on consumption of a single final good (c_{ot}) and non-wage compensating differentials for that occupation (b_{ot}):

$$u_{ot} = \ln b_{ot} + \ln c_{ot}. \tag{1}$$

The corresponding indirect utility function is:

$$u_{ot} = \ln b_{ot} + \ln w_{ot} - \ln p_t, \tag{2}$$

where w_{ot} denotes the wage, which can differ across occupations because of mobility costs; p_t is the unit expenditure function dual to the goods consumption index; we choose the consumption good price index as our numeraire, such that $p_t = 1$.

After supplying labor and spending wage income on consumption in each period t, workers observe idiosyncratic mobility shocks (ϵ_{ot}) and decide which occupation to move to for the next period $t+1$. The value function for a worker in occupation o in period t ($\mathbb{V}^w_{ot}$) is equal to the current flow of utility in that occupation plus the expected continuation value from the optimal choice of occupation:

$$\mathbb{V}^w_{ot} = u_{ot} + \max_{\{i \in \mathbb{O}\}} \left\{ \beta \mathbb{V}^w_{it+1} - \ln \kappa_{iot} + \rho \epsilon_{it} \right\}, \tag{3}$$

where β is the discount rate and κ_{iot} capture bilateral mobility frictions between occupations; we assume that these bilateral mobility frictions satisfy $\kappa_{oot} = 1$ and $\kappa_{iot} > 1$ for $i \neq o$.

We assume log utility for workers, because they live hand-to-mouth, and hence there is no role for the intertemporal elasticity of substitution in their consumption decisions. We make the conventional assumption that these idiosyncratic mobility shocks are drawn from an extreme value distribution with CDF $F(\epsilon) = e^{-e^{(-\epsilon - \overline{\gamma})}}$, where $\overline{\gamma}$ is the Euler–Mascheroni constant; the parameter ρ controls the dispersion of these idiosyncratic mobility shocks.

2.2 Capitalist's Problem

At the beginning of period t, the representative capitalist inherits a stock of wealth a_t, which is stored in occupation-specific capital $\{k_{ot}\}$. The representative capitalist chooses consumption and how much investment to allocate to each occupation (and hence overall wealth accumulation) . We denote the return to saving by R_t and the end of period price per unit of capital stored in occupation o by q_{ot-1}. Wealth (a_t) satisfies:

$$a_t = \sum_{o \in \mathbb{O}} q_{ot-1} k_{ot}. \tag{4}$$

We assume that investment in each occupation is subject to adjustment costs. We model these adjustment costs following Lucas and Prescott (1971), Hercowitz and Sampson (1991) and Liu and Ma (2023). If the capitalist invests χ_{ot} units of the consumption index in occupation o in time t, the stock of capital in occupation o at time $t+1$ is given by:

$$k_{ot+1} = \eta_{ot} k_{ot}^{1-\delta} \chi_{ot}^{\delta}, \tag{5}$$

where δ controls the magnitude of adjustment costs and η_{ot} captures the efficiency of the investment technology for each occupation.

In this specification, adjustment costs arise because of diminishing marginal returns to investing additional units of final goods output (χ_{ot}) to accumulate more capital (k_{ot+1}), given the existing stock of capital (k_{ot}). As a result of these diminishing returns, concentrating investment in an occupation in a single time period generates less future capital in that occupation than spreading investment out across time periods. We assume that capital is rented in a competitive capital market in each occupation at rental rate r_{ot}. No-arbitrage implies that the return to investment (incorporating adjustment costs) is equalized across occupations. But the rental rate (r_{ot}) can differ across occupations because of the adjustment costs.

We assume that the representative capitalist has logarithmic intertemporal utility. She chooses consumption and investment to maximize intertemporal utility subject to her budget constraint:

$$\max_{\{c_t, a_{t+1}\}} \sum_{s=0}^{\infty} \beta^{t+s} \ln c_{t+s}^k, \tag{6}$$
$$\text{subject to} \qquad c_t^k + a_{t+1} = R_t a_t,$$

where we use the superscript k to denote the value of variables for the representative capitalist, such that c_t^k is her consumption of the single final good; and β is the discount rate.

2.3 Final Goods Production Technology

Markets are perfectly competitive. We assumed a nested constant elasticity of substitution (CES) production structure. In the upper nest, the final good is produced by performing a set of procedures $o \in \mathbb{O}_{jt}$, each of which corresponds to an occupation. In the lower nest for each procedure, tasks are either undertaken by labor or are automated and undertaken by capital.

In the upper nest, cost minimization and zero profits imply that the price of the final good (p_t) equals its unit cost of production:

$$p_t = \left[\sum_{o \in \mathbb{O}} \left(\frac{q_{ot}}{z_{ot}} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \tag{7}$$

where q_{ot} denotes the cost of performing each procedure; z_{ot} captures productivity for each occupation; we allow productivity to evolve unevenly across occupations over time, such that z_{ot} measures non-neutral technical change; more broadly, z_{ot} shifts the demand for procedures conditional on price, and hence encompasses changes in both demand and productivity as sources of structural transformation; σ is the elasticity of substitution across occupations.

Applying Shephard's Lemma to equation (7), the share of each procedure in unit costs in the upper nest depends on relative costs and productivity across occupations:

$$s_{ot} = \frac{\left(q_{ot}/z_{ot}\right)^{1-\sigma}}{\sum_{m \in \mathbb{O}} \left(q_{mt}/z_{mt}\right)^{1-\sigma}}. \tag{8}$$

Cost minimization and zero profits imply that the cost of performing each procedure o (q_{ot}) in the lower nest is a constant elasticity function of the wage (w_{ot}) and rental rate (r_{ot}):

$$q_{ot} = \left[(1 - \alpha_{ot})\, w_{ot}^{1-\zeta} + \alpha_{ot} r_{ot}^{1-\zeta} \right]^{\frac{1}{1-\zeta}}, \tag{9}$$

where ζ is the elasticity of substitution between labor and capital; for simplicity, we assume that a single elasticity of substitution for all procedures, but it is straightforward to allow this elasticity to vary across procedures.

The parameter α_{ot} controls the relative weight of capital and labor in unit costs and captures the automation of tasks within each procedure. We derive this CES unit cost function from a microfoundation with a unit continuum of tasks within each procedure in Online Appendix B.3. Each task can be performed by either labor or automated capital. Each factor of production is subject to an idiosyncratic shock to its productivity for each task. Labor and capital are perfect substitutes at the level of individual tasks. Therefore, each task is undertaken by the factor that can perform it at the lowest cost. If the idiosyncratic productivity shocks have an extreme value function, we show that the unit cost function across the continuum of tasks takes the CES form in equation (9). The weighting parameter (α_{ot}) depends on the scale parameter of the extreme value distribution and captures the fraction of tasks that have been automated.

Applying Shephard's Lemma to equation (9), the share of labor in unit costs is:

$$f_{ot} = \frac{(1 - \alpha_{ot})\, w_{ot}^{1-\zeta}}{(1 - \alpha_{ot})\, w_{ot}^{1-\zeta} + \alpha_{ot} r_{ot}^{1-\zeta}}. \tag{10}$$

Increases in α_{ot} capture automation shocks (e.g., from computerization) that raise the productivity of capital relative to labor in performing tasks within a procedure, thereby increasing the fraction of tasks undertaken by capital and the share of capital in unit costs.

2.4 Market Clearing

Market clearing requires that the value of production for each procedure is equal to the value of expenditure on that procedure:

$$w_{ot}\ell_{ot} + r_{ot}k_{ot} = s_{ot} \sum_{m \in \mathbb{O}} \left[w_{mt}\ell_{mt} + r_{mt}k_{mt} \right], \tag{11}$$

where zero profits implies that the value of production for each procedure is equal to the value of factor payments.

2.5 General Equilibrium

Given the state variables $\{\ell_{ot}, k_{ot}\}$, the general equilibrium of the economy is the set of allocations and prices such that firms choose factor inputs to maximize profits; workers make consumption

and occupation mobility decisions to maximize utility; capitalists choose investment decisions to maximize their value functions; and prices clear all markets.

The evolution of the endogenous variables is driven by five fundamentals: (i) Automation (α_{ot}); (ii) Productivity (z_{ot}); (iii) Non-wage compensating differentials (b_{ot}); (iv) Investment efficiency (η_{ot}); (v) worker mobility frictions (κ_{oit}). We now summarize the three blocks of general equilibrium conditions that determine the evolution of the endogenous variables in response to changes in these five fundamentals: (i) A static block; (ii) A worker dynamic block; (iii) A capitalist dynamic block.

2.5.1 Static Block of Equations

The static block of equations determines the equilibrium values of wages (w_{ot}), rental rates (r_{ot}), procedure cost shares (s_{ot}), labor shares (f_{ot}), and procedure prices (q_{ot}) given the state variables $\{\ell_{ot}, k_{ot}\}$ and a choice of numeraire ($p_t = 1$):

$$p_t = \left[\sum_{o\in\mathbb{O}} \left(\frac{q_{ot}}{z_{ot}}\right)^{1-\sigma}\right] = 1, \tag{12}$$

$$s_{ot} = \frac{\left(q_{ot}/z_{ot}\right)^{1-\sigma}}{\sum_{m\in\mathbb{O}}\left(q_{mt}/z_{mt}\right)^{1-\sigma}}, \tag{13}$$

$$q_{ot} = \left[\left(1-\alpha_{ot}\right) w_{ot}^{1-\zeta} + \alpha_{ot} r_{ot}^{1-\zeta}\right]^{\frac{1}{1-\zeta}}, \tag{14}$$

$$f_{ot} = \frac{\left(1-\alpha_{ot}\right) w_{ot}^{1-\zeta}}{\left(1-\alpha_{ot}\right) w_{ot}^{1-\zeta} + \alpha_{ot} r_{ot}^{1-\zeta}}, \tag{15}$$

$$w_{ot}\ell_{ot} + r_{ot}k_{ot} = s_{ot}\sum_{m\in\mathbb{O}}\left(w_{mt}\ell_{mt} + r_{mt}k_{mt}\right). \tag{16}$$

2.5.2 Worker Dynamic Block of Equations

The worker dynamic block of equations determines the evolution of value functions (v_{ot}^w) and the labor state variables (ℓ_{ot}). Using the properties of the extreme value distribution, the solution for the expected value of workers in occupation o in period t from equation (3) is:

$$v_{ot}^w = u_{ot} + \rho \ln \sum_{i\in\mathbb{O}} \left(\exp\left(\beta v_{it+1}^w\right)/\kappa_{iot}\right)^{1/\rho}, \tag{17}$$

and the corresponding evolution of employment across occupations satisfies:

$$\ell_{ot+1} = \sum_{i\in\mathbb{O}} D_{iot}\ell_{it}, \qquad D_{iot} \equiv \frac{\left(\exp\left(\beta v_{ot+1}^w\right)/\kappa_{oit}\right)^{1/\rho}}{\sum_{m\in\mathbb{O}}\left(\exp\left(\beta v_{mt+1}^w\right)/\kappa_{mit}\right)^{1/\rho}}, \tag{18}$$

where we define D_{iot} as the probability of outmigrating from occupation i to occupation o at time t, and define an analogous probability of inmigrating to occupation o from occupation i ($E_{oit} \equiv D_{iot}\ell_{it}/\ell_{ot+1}$).

From this worker dynamic block, idiosyncratic mobility shocks (whose dispersion is controlled by ρ) and bilateral mobility costs (captured by κ_{oit}) imply that reallocations of workers across occupations in response to shocks to fundamentals occur gradually over time.

2.5.3 Capitalist Dynamic Block of Equations

The capitalist dynamic block of equations determines the evolution of the capital state variables (k_{ot}). We begin by characterizing the representative capitalist's consumption-investment decision, before characterizing her allocation of investment across occupations. With logarithmic intertemporal utility, the optimal consumption-investment decisions that solve (6) are characterized by a constant equilibrium investment rate, as in Moll (2014):

$$\begin{aligned} c_t &= (1-\beta)\, R_t a_t, \\ a_{t+1} &= \beta R_t a_t. \end{aligned} \tag{19}$$

The initial stock of wealth is optimally allocated across occupations to equate the marginal return to investment with the marginal cost of investment in terms of a unit of the final good:

$$q_{ot}\frac{\partial\left(\eta_{ot}k_{ot}^{1-\delta}\chi_{ot}^{\delta}\right)}{\partial\chi_{ot}} = q_{ot}\delta\eta_{ot}\left(\frac{k_{ot}}{\chi_{ot}}\right)^{1-\delta} = 1. \tag{20}$$

Using this first-order condition in the investment technology (5), the common return to investment across all occupations satisfies:

$$R_t = \frac{r_{ot} + q_{ot}\frac{\partial\left(\eta_{ot}k_{ot}^{1-\delta}\chi_{ot}^{\delta}\right)}{\partial k_{ot}}}{q_{ot-1}} = \frac{r_{ot} + q_{ot}\left(1-\delta\right)\eta_{ot}\left(\chi_{ot}/k_{ot}\right)^{\delta}}{q_{ot-1}}, \tag{21}$$

which makes clear that the rental rates (r_{ot}) differ across occupations along the transition path to steady-state, because of the adjustment costs.

Using this common return on investment together with market clearing for the final good and the representative capitalist's intertemporal budget constraint, the equilibrium allocation of investment across occupations satisfies:

$$\frac{\chi_{ot-1}}{\sum_{o\in\mathbb{O}}\chi_{ot-1}} = (1-\beta+\beta\delta)\sum_{s=0}^{\infty}\left(\beta\left(1-\delta\right)\right)^{s}\frac{r_{ot+s}k_{ot+s}}{\sum_{o\in\mathbb{O}}r_{ot+s}k_{ot+s}}, \tag{22}$$

which has the intuitive interpretation that relative investments across occupations depend on discounted expected relative future returns to capital. Using this equilibrium allocation in the

investment technology (5), the implied law of motion for capital is:

$$\ln k_{ot+1} = \left[\begin{array}{c} \ln \eta_{ot} + (1-\delta) \ln k_{ot} \\ +\delta \ln \left[\beta\delta \left(\sum_{m \in \mathbb{O}} r_{mt} k_{mt} \right) \left(\sum_{s=1}^{\infty} \left(\beta \left(1-\delta \right) \right)^{s} \frac{r_{ot+s} k_{ot+s}}{\sum_{m \in \mathbb{O}} r_{mt+s} k_{mt+s}} \right) \right] \end{array} \right]. \tag{23}$$

From this capitalist dynamic block, adjustment costs (captured by δ) imply that reallocations of capital across procedures in response to shocks to fundamentals also occur gradually over time.

3 Employment and Wage Dynamics

The evolution of employment and wages in response to automation (α_{ot}), non-neutral technical change (z_{ot}) and other shocks to fundamentals is driven by the interaction between labor and capital dynamics. We now characterize this interaction between these two sources of dynamics. We begin by showing that we can use dynamic exact-hat algebra to solve for the economy's transition path in response to shocks to fundamentals, using only the observed endogenous variables in the data, and without requiring information on the unobserved levels of fundamentals. To provide a further analytical characterization of the model's transition dynamics, we next linearize the model to obtain a closed-form solution for the economy's transition path.

3.1 Dynamic Exact-Hat Algebra

We suppose that we observe the economy somewhere along the transition path towards an unobserved steady state. Given the initial observed endogenous variables of the model, we show that we are able to solve for the economy's transition path in time differences ($\dot{x}_{ot+1} = x_{ot+1}/x_{ot}$) for any anticipated convergent sequence of future changes in fundamentals, without having to solve for the initial level of fundamentals.

Proposition 1. ***Dynamic Exact Hat Algebra.*** *Given an initial observed allocation of the economy* ($\{\ell_{o0}\}_{o=1}^{O}$, $\{k_{o0}\}_{o=1}^{O}$, $\{k_{o1}\}_{o=1}^{O}$, $(\{D_{io0}\}_{i,o=1}^{O})$*, and a convergent sequence of future changes in fundamentals under perfect foresight:*

$$\left\{ \{\dot{\alpha}_{ot}\}_{o=1}^{O}, \{\dot{z}_{ot}\}_{o=1}^{O}, \left\{\dot{b}_{ot}\right\}_{o=1}^{O}, \{\dot{\eta}_{ot}\}_{o=1}^{O}, \{\dot{\kappa}_{oit}\}_{o,i=1}^{O} \right\}_{t=1}^{\infty},$$

the solution for the sequence of changes in the model's endogenous variables does not require information on the level of fundamentals:

$$\left\{ \{\alpha_{ot}\}_{i,j=1}^{N}, \{z_{ot}\}_{o=1}^{O}, \{b_{ot}\}_{o=1}^{O}, \{\eta_{ot}\}_{i,j=1}^{N}, \{\kappa_{oit}\}_{o,i=1}^{O} \right\}_{t=1}^{\infty}.$$

Proof. See Online Appendix E. □

Intuitively, we use the initial observed endogenous variables and the equilibrium conditions of the model to control for the unobserved initial level of fundamentals. We can use this proposition and the observed data for any year to solve for the unobserved steady-state in the absence of any further changes in fundamentals for that year. Therefore, we can examine the extent to which the economy has fallen further behind or closed the gap to steady-state over time, as fundamental shocks have occurred during our sample period.

Additionally, we can use this proposition and the observed data to solve for counterfactuals for the economy's transition path in response to assumed sequences of future changes in fundamentals. We can implement this analysis starting in any year during our sample. Therefore, we can examine the role of automation, non-neutral technical change, other shocks to fundamentals, and initial conditions in understanding the evolution of the observed endogenous variables over time.

Finally, we can also use these dynamic exact-hat counterfactuals to predict the future impact of emergent new technologies, such as artificial intelligence. Importantly, we are not required to assume that the economy is in steady-state, and can implement all of this analysis for an economy that begins anywhere along the transition path to an unobserved steady-state.

3.2 Model Inversion

Our dynamic exact-hat algebra results in Proposition 1 allow us to undertake counterfactuals for future changes in fundamentals, without needing to know the initial levels of fundamentals. We show that we can invert the equilibrium conditions of the model to recover the unobserved fundamentals (α_{ot}, z_{ot}, b_{ot}, η_{ot}, κ_{not}), given the observed data on employment (ℓ_{ot}), wages (w_{ot}), capital stocks (k_{ot}), capital payments ($r_{ot}k_{ot}$) and worker mobility (D_{oit}) and the model parameters. We use these solutions for the unobserved fundamentals to compute measures of automation ($\alpha_{ot}/\alpha_{ot-T}$), non-neutral technical change (z_{ot}/z_{ot-T}) and other shocks to fundamentals over time. Using these measures of shocks to fundamentals in our dynamic exact-hat algebra counterfactuals in the previous subsection, we can quantify the role of convergence to the initial steady-state, automation, non-neutral technical change and other fundamental shocks in understanding the observed changes in employment and wages in the data.

Proposition 2. ***Model Inversion.*** *Suppose that we observe sequences of values of the models endogenous variables: ($\{\ell_{ot}\}_{t=0}^{T}$, $\{w_{ot}\}_{t=0}^{T}$, $\{k_{ot}\}_{0=1}^{T}$, $\{r_{ot}k_{ot}\}_{t=0}^{T}$, $\{D_{oit}\}_{t=0}^{T}$). From these observed sequences and given model parameters (β, ρ, δ, σ, ζ), we can recover unique sequences of fundamentals ($\{\alpha_{ot}\}_{t=1}^{T}$,$\{z_{ot}\}_{t=1}^{T}$,$\{b_{ot}\}_{t=1}^{T}$, $\{\eta_{ot}\}_{t=1}^{T}$, $\{\kappa_{oit}\}_{t=1}^{T}$) for which the observed data are consistent with an equilibrium of the model.*

Proof. Online Appendix F. □

Proposition 2 assumes that we observe the economy in an initial year somewhere along the transition path towards an unobserved steady-state. From that point in time onwards, the evolution of the endogenous variables reflects a combination of convergence towards that unobserved initial steady-state, and impact of shocks to fundamentals that occur during the sample period. Given the observed sequence of endogenous variables in the data, we can recover unique sequences of shocks to fundamentals that occur during the sample period for which the evolution of the endogenous variables is an equilibrium of the model.

This proposition does not make any assumptions about how far the economy is from steady-state in the initial year. It also does not impose restrictions on agents' unobserved expectations about the future trajectory of fundamentals (except for our assumption of perfect foresight). Instead we use the observed worker mobility flows and investment decisions to control for these unobserved expectations. We require one more year of data for the endogenous variables (from year 0 to T) than the period for which we can recover fundamentals (from year 1 to T), because we need at least two periods for worker mobility flows and capital stocks, in order to use observed worker mobility flows and investment to control for agents' unobserved expectations for the future trajectory of fundamentals.

Proposition 2 has an important corollary. If we start from the observed value of the endogenous variables in the initial year, and undertake a dynamic exact-hat algebra counterfactual feeding in the future sequences of all fundamentals recovered from the model inversion, we exactly recover the sequences of the endogenous variables observed in the data. We use this corollary to explore the role of automation shocks, changes in other fundamentals, and initial conditions in explaining the employment and wage polarization and changes in wage inequality observed in the data. We undertake dynamic exact-hat algebra counterfactuals changing no fundamentals (such that only convergence based on initial conditions occurs) and then progressively adding additional fundamental shocks (e.g., automation, or automation plus non-neutral technical change).

Comparing the observed data (which include the effects of all fundamental shocks and initial conditions) to these counterfactuals in which only some of these components change reveals the relative importance of initial conditions and different fundamentals in understanding the employment and wage polarization and changes in wage inequality observed in the data.

3.3 Linearization

We now linearize the model to characterize analytically the model's transition dynamics. We suppose that we observe the capital state variable ($\boldsymbol{k_t}$) for two time periods $t = 0$ and $t = 1$, the employment state variable ($\boldsymbol{\ell}_t$) for $t = 0$, wages ($\boldsymbol{w_t}$) and rental rates ($\boldsymbol{r_t}$) for $t = 0$, and the worker mobility matrix ($\boldsymbol{D_t}$) for two time periods $t = 0$ and $t = 1$. The economy need not be in steady state at $t = 0$, but we assume that it is on a convergence path towards a steady state

with constant fundamentals ($\boldsymbol{\alpha}$, $\boldsymbol{z}$, $\boldsymbol{b}$, $\boldsymbol{\eta}$, $\boldsymbol{\kappa}$). We refer to the steady state implied by these initial fundamentals as the *initial* steady state. We use a tilde above a variable to denote a log deviation from this initial steady state (e.g., $\widetilde{\ell}_{ot+1} = \ln \ell_{ot+1} - \ln \ell_o^*$).

We begin by totally differentiating the conditions for general equilibrium around this unobserved initial steady state, holding constant countries' labor endowments. We thus obtain a system of linear equations that fully characterizes the economy's transition path up to first-order, as reported in Online Appendix C. We next show that this system of linearized equations can be reduced to a second-order difference equation in the employment and capital state variables ($\widetilde{\boldsymbol{\ell}}_{\boldsymbol{t}}$, $\widetilde{\boldsymbol{k}}_{\boldsymbol{t}}$) and changes to fundamentals. For expositional convenience, we focus here on the simplest form of changes in fundamentals, such that agents at time $t = 0$ learn about a one-time permanent shock to fundamentals from time $t = 1$ onwards. However, analogous results hold for any expected convergent sequence of future shocks to fundamentals under perfect foresight.

We define measures of incoming and outgoing shocks to worker mobility frictions, which aggregate bilateral changes in these frictions using outmigration (D_{iot}) and inmigration (E_{oit}) shares: $\widetilde{\kappa}_{it}^{out} \equiv \sum_{o=1}^{O} D_{iot}\widetilde{\kappa}_{oit}$ and $\widetilde{\kappa}_{ot}^{in} \equiv \sum_{i=1}^{O} E_{oit}\widetilde{\kappa}_{oit}$.

Proposition 3. ***State Variables.*** *Suppose that the economy at time $t = 0$ is on a convergence path toward an initial steady state with constant fundamentals ($\boldsymbol{\alpha}$, $\boldsymbol{z}$, $\boldsymbol{b}$, $\boldsymbol{\eta}$, $\boldsymbol{\kappa}$). At time $t = 0$, agents learn about one-time, permanent shocks to fundamentals ($\widetilde{\boldsymbol{f}} \equiv \begin{bmatrix} \widetilde{\boldsymbol{\alpha}} & \widetilde{\boldsymbol{z}} & \widetilde{\boldsymbol{b}} & \widetilde{\boldsymbol{\eta}} & \widetilde{\boldsymbol{\kappa}}^{in} & \widetilde{\boldsymbol{\kappa}}^{out} \end{bmatrix}'$) from time $t = 1$ onwards. The evolution of the economy's state variables $\widetilde{\boldsymbol{x}}_{\boldsymbol{t}} = \begin{bmatrix} \widetilde{\boldsymbol{\ell}}_{\boldsymbol{t}} \\ \widetilde{\boldsymbol{k}}_{\boldsymbol{t}} \end{bmatrix}$ from time $t = 1$ onwards satisfies the following second-order difference equation:*

$$\boldsymbol{\Psi}\widetilde{\boldsymbol{x}}_{t+2} = \boldsymbol{\Gamma}\widetilde{\boldsymbol{x}}_{t+1} + \boldsymbol{\Theta}\widetilde{\boldsymbol{x}}_{t} + \boldsymbol{\Pi}\widetilde{\boldsymbol{f}}, \tag{24}$$

where the matrices ($\boldsymbol{\Psi}$, $\boldsymbol{\Gamma}$, $\boldsymbol{\Theta}$, $\boldsymbol{\Pi}$) are functions of the mobility matrices ($\boldsymbol{D}$, $\boldsymbol{E}$) and model parameters (β, ρ, δ, σ, ζ), as defined in Online Appendix D.4.

Proof. See Online Appendix D.4. □

We solve this matrix system of equations using the method of undetermined coefficients following Uhlig (1999) to obtain a closed-form solution for the evolution of the state variables $\{\widetilde{\boldsymbol{x}}_{\boldsymbol{t}}\}_{t=1}^{\infty}$ in terms of an impact matrix ($\boldsymbol{R}$), which captures the initial impact of the fundamental shocks, and a transition matrix ($\boldsymbol{P}$), which governs the updating of the state variables over time.

Proposition 4. ***Transition Matrix.*** *Suppose that the economy at time $t = 0$ is on a convergence path toward an initial steady state with constant fundamentals ($\boldsymbol{\alpha}$, $\boldsymbol{z}$, $\boldsymbol{b}$, $\boldsymbol{\eta}$, $\boldsymbol{\kappa}$). At time $t = 0$, agents learn about one-time, permanent shocks to fundamentals ($\widetilde{\boldsymbol{f}} \equiv \begin{bmatrix} \widetilde{\boldsymbol{\alpha}} & \widetilde{\boldsymbol{z}} & \widetilde{\boldsymbol{b}} & \widetilde{\boldsymbol{\eta}} & \widetilde{\boldsymbol{\kappa}}^{in} & \widetilde{\boldsymbol{\kappa}}^{out} \end{bmatrix}'$) from*

time $t = 1$ onwards. There exists a $2O \times 2O$ transition matrix ($\boldsymbol{P}$) and a $2O \times 5O$ impact matrix ($\boldsymbol{R}$) such that the second-order difference equation system in (24) has the closed-form solution:

$$\widetilde{\boldsymbol{x}}_t = \boldsymbol{P}\widetilde{\boldsymbol{x}}_{t-1} + \boldsymbol{R}\widetilde{\boldsymbol{f}}. \tag{25}$$

The transition matrix $\boldsymbol{P}$ satisfies:

$$\boldsymbol{P} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{-1},$$

where $\boldsymbol{\Lambda}$ is a diagonal matrix of $2O$ stable eigenvalues $\{\lambda_o\}_{o=1}^{2O}$ and $\boldsymbol{U}$ is a matrix stacking the corresponding O eigenvectors $\{\boldsymbol{u}_k\}_{k=1}^{N}$. The impact matrix ($\boldsymbol{R}$) is given by:

$$\boldsymbol{R} = \left(\boldsymbol{\Psi}\boldsymbol{P} + \boldsymbol{\Psi} - \boldsymbol{\Gamma}\right)^{-1}\boldsymbol{\Pi},$$

where ($\boldsymbol{\Psi}$, $\boldsymbol{\Gamma}$, $\boldsymbol{\Theta}$, $\boldsymbol{\Pi}$) are the matrices from the system of second-order difference equations (24).

Proof. See Online Appendix D.5. □

The solutions for these impact and transition matrices ($\boldsymbol{R}$, $\boldsymbol{P}$) depend only on the mobility matrices ($\boldsymbol{D}$, $\boldsymbol{E}$) and parameters (β, ρ, δ, σ, ζ). Given this closed-form solution for the state variables $\{\widetilde{\boldsymbol{x}}_{\boldsymbol{t}}\}_{t=1}^{\infty}$, we can recover all other endogenous variables as linear functions of these state variables, as shown in Online Appendix D.3.

3.4 Convergence Dynamics Versus Fundamental Shocks

Using Proposition 4, the transition path of the economy's state variables can be additively decomposed into the contributions of convergence dynamics given initial conditions and fundamental shocks. Applying equation (25) across time periods, we obtain:

$$\ln \boldsymbol{x}_t - \ln \boldsymbol{x}_{-1} = \underbrace{\sum_{s=0}^{t} \boldsymbol{P}^s \left(\ln \boldsymbol{x}_0 - \ln \boldsymbol{x}_{-1}\right)}_{\substack{\text{convergence given} \\ \text{initial fundamentals}}} + \underbrace{\sum_{s=0}^{t-1} \boldsymbol{P}^s \boldsymbol{R}\widetilde{\boldsymbol{f}}}_{\substack{\text{dynamics from} \\ \text{fundamental shocks}}} \quad \text{for all } t \geq 1. \tag{26}$$

In the absence of shocks to fundamentals ($\widetilde{\boldsymbol{f}} = \boldsymbol{0}$), the second term on the right-hand side of equation (26) is zero. In this case, the evolution of the state variables is shaped solely by convergence dynamics given initial conditions, and converges over time to:

$$\ln \boldsymbol{x}^*_{\text{initial}} = \lim_{t\to\infty} \ln \boldsymbol{x}_t = \ln \boldsymbol{x}_{-1} + \left(\boldsymbol{I} - \boldsymbol{P}\right)^{-1} \left(\ln \boldsymbol{x}_0 - \ln \boldsymbol{x}_{-1}\right), \tag{27}$$

where $\left(\boldsymbol{I} - \boldsymbol{P}\right)^{-1} = \sum_{s=0}^{\infty} \boldsymbol{R}^s$ is well-defined under the condition that the spectral radius of $\boldsymbol{P}$ is smaller than one.

In contrast, if the economy is initially in a steady state at time $t = 0$, the first term on the right-hand side of equation (26) is zero. In this case, the transition path of the state variables is solely driven by the second term for fundamental shocks, and follows:

$$\widetilde{\boldsymbol{x}}_t = \ln \boldsymbol{x}_t - \ln \boldsymbol{x}_0 = \sum_{s=0}^{t-1} \boldsymbol{P}^s \boldsymbol{R}\widetilde{\boldsymbol{f}} = \left(\boldsymbol{I} - \boldsymbol{P}^t\right)\left(\boldsymbol{I} - \boldsymbol{P}\right)^{-1} \boldsymbol{R}\widetilde{\boldsymbol{f}} \quad \text{for all } t \geq 1. \tag{28}$$

In period $t = 1$ when the shocks occur, the response of the state variables is $\widetilde{\boldsymbol{x}}_1 = \boldsymbol{R}\widetilde{\boldsymbol{f}}$. Taking the limit as $t \to \infty$ in equation (28), the comparative steady-state response is

$$\lim_{t\to\infty} \widetilde{\boldsymbol{x}}_t = \ln \boldsymbol{x}^*_{\text{new}} - \ln \boldsymbol{x}^*_{\text{initial}} = \left(\boldsymbol{I} - \boldsymbol{P}\right)^{-1} \boldsymbol{R}\widetilde{\boldsymbol{f}}. \tag{29}$$

A key implication of this additive separability in equation (26) is that we can examine the economy's dynamic response to fundamental shocks separately from its convergence towards an initial steady state with unchanged fundamentals. Therefore, without loss of generality, we focus in the remainder of this section on an economy that is initially in steady state.

3.5 Spectral Analysis of the Transition Matrix $\boldsymbol{P}$

We now provide a further analytical characterization of the economy's dynamic response to fundamental shocks using a spectral analysis of the transition matrix. We show that the speed of convergence to steady state and the evolution of the state variables along the transition path can be written solely in terms of the eigenvalues and eigenvectors of this transition matrix. Since all other endogenous variables are linear functions of the state variables, we also obtain impulse responses of each of the endogenous variables with respect to shocks to fundamentals.

3.5.1 Eigendecomposition of the Transition Matrix

We begin by using the eigendecomposition of the transition matrix, $\boldsymbol{P} \equiv \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{V}$, where $\boldsymbol{\Lambda}$ is a diagonal matrix of eigenvalues arranged in decreasing order by absolute values, and $\boldsymbol{V} = \boldsymbol{U}^{-1}$. For each eigenvalue λ_h, the h-th column of $\boldsymbol{U}$ ($\boldsymbol{u_h}$) and the h-th row of $\boldsymbol{V}$ ($\boldsymbol{v_h'}$) are the corresponding right- and left-eigenvectors of $\boldsymbol{P}$, respectively, such that

$$\lambda_h \boldsymbol{u}_h = \boldsymbol{P}\boldsymbol{u}_h, \qquad \lambda_h \boldsymbol{v}_h' = \boldsymbol{v}_h'\boldsymbol{P}.$$

That is, $\boldsymbol{u}_h$ ($\boldsymbol{v}_h'$) is the vector that, when left-multiplied (right-multiplied) by $\boldsymbol{P}$, is proportional to itself but scaled by the corresponding eigenvalue λ_h.[1] We refer to $\boldsymbol{u}_h$ simply as eigenvectors. Both $\{\boldsymbol{u}_h\}$ and $\{\boldsymbol{v}_h'\}$ are bases that span the $2O$-dimensional state space.

[1]Note that $\boldsymbol{P}$ need not be symmetric. This eigendecomposition exists if the transition matrix has distinct eigenvalues. We construct the right-eigenvectors such that the 2-norm of $\boldsymbol{u}_h$ is equal to 1 for all h, where note that $\boldsymbol{v}_i'\boldsymbol{u}_h = 1$ for $i = h$ and $\boldsymbol{v}_i'\boldsymbol{u}_h = 0$ otherwise.

3.5.2 Eigen-shock

We next introduce a particular type of shock to fundamentals that proves useful for characterizing the model's transition dynamics. We define an *eigen-shock* as a shock to fundamentals ($\widetilde{\boldsymbol{f}}_{(h)}$) for which the initial impact of the shock on the state variables ($\boldsymbol{R}\widetilde{\boldsymbol{f}}_{(h)}$) coincides with a real eigenvector of the transition matrix ($\boldsymbol{u}_h$). With $2 \times O$ state variables ($\widetilde{\boldsymbol{x}}$) and $5 \times O$ fundamental shocks ($\widetilde{\boldsymbol{\alpha}}$, $\widetilde{\boldsymbol{z}}$, $\widetilde{\boldsymbol{b}}$, $\widetilde{\boldsymbol{\eta}}$, $\widetilde{\boldsymbol{\kappa}}^{in}$, $\widetilde{\boldsymbol{\kappa}}^{out}$), the space of fundamental shocks is of higher dimension than the space of state variables. Therefore, many fundamental shocks generate identical time paths for the state variables. In fact, for any fundamental shock vector ($\widetilde{\boldsymbol{f}}$), there exist automation and productivity shock vectors ($\widetilde{\boldsymbol{\alpha}}$, $\widetilde{\boldsymbol{z}}$) that generate the same time path of the state variables. For expositional simplicity, we define the eigen-shocks in terms of shocks to automation and productivity ($\widetilde{\boldsymbol{\alpha}}$, $\widetilde{\boldsymbol{z}}$), setting all other shocks equal to zero.[2] Consequentially, the impact of the eigen-shocks $\left\{ \widetilde{\boldsymbol{f}}_{(h)} \right\}_{o=1}^{2O}$ form a basis that spans the $2O$-dimensional state space. Each eigenvector of $\boldsymbol{P}$ ($\boldsymbol{u}_h$) has a corresponding eigen-shock for which $\boldsymbol{R}\widetilde{\boldsymbol{f}}_{(h)} = \boldsymbol{u}_h$.

In general, there is no reason why any vector of empirical shocks to fundamentals across countries should correspond to an eigen-shock. But we can use these eigen-shocks to characterize the impact of any empirical shock using the following two properties. First, we can solve for these eigen-shocks from the observed data, because the impact matrix ($\boldsymbol{R}$) and the transition matrix ($\boldsymbol{P}$) depend solely on our observed mobility matrices ($\boldsymbol{D}$, $\boldsymbol{E}$) and the structural parameters of the model (β, ρ, δ, σ, ζ). Second, the initial and dynamic impact on the state variables from any vector of empirical shocks to fundamentals ($\widetilde{\boldsymbol{f}}$) can be equivalently expressed as a linear combination of the impact from eigen-shocks ($\widetilde{\boldsymbol{f}}_{(h)}$), where the weights or loadings in this linear combination can be recovered from a linear projection (regression) of the initial impact from the observed shocks ($\boldsymbol{R}\widetilde{\boldsymbol{f}}$) on the initial impact from the eigen-shocks ($\boldsymbol{R}\widetilde{\boldsymbol{f}}_{(h)}$). Using this property, the transition path of the state variables in response to any vector of empirical shocks to fundamentals can be expressed solely in terms of the eigenvalues and eigenvectors of the transition matrix, as summarized in the following proposition.

Proposition 5. ***Spectral Analysis.*** *Consider an economy that is initially in steady state at time $t = 0$ when agents learn about one-time, permanent shocks to fundamentals ($\widetilde{\boldsymbol{f}} \equiv \begin{bmatrix} \widetilde{\boldsymbol{\alpha}} & \widetilde{\boldsymbol{z}} & \widetilde{\boldsymbol{b}} & \widetilde{\boldsymbol{\eta}} & \widetilde{\boldsymbol{\kappa}}^{in} & \widetilde{\boldsymbol{\kappa}}^{out} \end{bmatrix}'$) from time $t = 1$ onwards. The transition path of the state variables ($\boldsymbol{x_t}$) can be written as a linear combination of the eigenvalues (λ_h) and eigenvectors ($\boldsymbol{u}_h$) of the*

[2] Recall from Proposition 3 that the dimension of the state space is $2 \times O$, whereas the dimension of the fundamental shocks is $5 \times O$, because $\widetilde{\boldsymbol{f}}$ includes shocks to automation, productivity, non-wage compensating differentials, and investment efficiency ($\widetilde{\boldsymbol{\alpha}}$, $\widetilde{\boldsymbol{z}}$, $\widetilde{\boldsymbol{b}}$, $\widetilde{\boldsymbol{\eta}}$) and aggregations of bilateral shocks to mobility frictions ($\widetilde{\boldsymbol{\kappa}}^{in}$, $\widetilde{\boldsymbol{\kappa}}^{out}$). Therefore, defining our eigen-shocks in terms of shocks to automation and productivity ($\widetilde{\boldsymbol{\alpha}}$, $\widetilde{\boldsymbol{z}}$) ensures that each eigenvector is associated with a unique eigen-shock (up to scale).

transition matrix:

$$\widetilde{\boldsymbol{x}}_t = \sum_{s=0}^{t-1} \boldsymbol{P}^s \boldsymbol{R}\widetilde{\boldsymbol{f}} = \sum_{h=1}^{2O} \frac{1-\lambda_h^t}{1-\lambda_h} \boldsymbol{u}_h \boldsymbol{v}_h' \boldsymbol{R}\widetilde{\boldsymbol{f}} = \sum_{h=2}^{2O} \frac{1-\lambda_h^t}{1-\lambda_h} \boldsymbol{u}_h \varrho_h, \tag{30}$$

where the weights in this linear combination (ϱ_h) can be recovered as the coefficients in a linear projection (regression) of the initial impact from the observed shocks ($\boldsymbol{R}\widetilde{\boldsymbol{f}}$) on the initial impact from the eigen-shocks ($\boldsymbol{R}\widetilde{\boldsymbol{f}}_{(h)}$).

Proof. The proposition follows from the eigendecomposition of the transition matrix: $\boldsymbol{P} \equiv \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{V}$, as shown in Online Appendix D.6. □

Another important property of an eigen-shock is that the speed of convergence to steady state, as measured by the half-life of convergence to steady state, depends solely on the associated eigenvalue of the transition matrix, as summarized in the following proposition.

Proposition 6. Speed of Convergence. *Consider an economy that is initially in steady state at time $t = 0$ when agents learn about one-time, permanent shocks to fundamentals ($\widetilde{\boldsymbol{f}} \equiv \begin{bmatrix} \widetilde{\boldsymbol{\alpha}} & \widetilde{\boldsymbol{z}} & \widetilde{\boldsymbol{b}} & \widetilde{\boldsymbol{\eta}} & \widetilde{\boldsymbol{\kappa}}^{in} & \widetilde{\boldsymbol{\kappa}}^{out} \end{bmatrix}'$) from time $t = 1$ onwards. Suppose that these shocks are an eigen-shock ($\widetilde{\boldsymbol{f}}_{(h)}$), for which the initial impact on the state variables at time $t = 1$ coincides with a real eigenvector ($\boldsymbol{u}_h$) of the transition matrix ($\boldsymbol{P}$): $\boldsymbol{R}\widetilde{\boldsymbol{f}}_{(h)} = \boldsymbol{u}_h$. The transition path of the state variables ($\boldsymbol{a_t}$) in response to such an eigen-shock ($\widetilde{\boldsymbol{f}}_{(h)}$) is :*

$$\widetilde{\boldsymbol{x}}_t = \sum_{j=2}^{2O} \frac{1-\lambda_j^t}{1-\lambda_j} \boldsymbol{u}_j \boldsymbol{v}_j' \boldsymbol{u}_h = \frac{1-\lambda_h^t}{1-\lambda_h} \boldsymbol{u}_h \quad \Longrightarrow \quad \ln \boldsymbol{x}_{t+1} - \ln \boldsymbol{x}_t = \lambda_h^t \boldsymbol{u}_h,$$

and the half-life of convergence to steady state is given by:

$$t_h^{(1/2)}\left(\widetilde{\boldsymbol{f}}\right) = -\left\lceil \frac{\ln 2}{\ln \lambda_h} \right\rceil,$$

for all state variables $h = 2, \cdots, 2N$, where $\widetilde{x}_{i\infty} = x^*_{i,new} - x^*_{i,initial}$, and $\lceil \cdot \rceil$ is the ceiling function.

Proof. The proposition follows from the eigendecomposition of the transition matrix ($\boldsymbol{P} \equiv \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{V}$), for the case of an eigen-shock in which the initial impact of the shocks to fundamentals on the state variables at time $t = 1$ coincides with a real eigenvector ($\boldsymbol{R}\widetilde{\boldsymbol{f}}_{(h)} = \boldsymbol{u}_h$) of the transition matrix ($\boldsymbol{P}$), as shown in Online Appendix D.7. □

From Proposition 6, the impact of an eigen-shock ($\widetilde{\boldsymbol{f}}_{(h)}$) on the state variables in each time period is always proportional to the corresponding eigenvector ($\boldsymbol{u}_h$), and decays exponentially at a rate determined by the associated eigenvalue (λ_h), as the economy converges to the new steady

state.[3] These eigenvalues fully summarize the economy's speed of convergence in response to eigen-shocks, even in our setting with many occupations and a rich mobility frictions.

In general, each eigen-shock ($\widetilde{\boldsymbol{f}}_{(h)}$) has a different speed of convergence (as captured by the associated eigenvalue λ_h), which reflects the fact that the speed of convergence to steady state does not only depend on the structural parameters of the model (β, ρ, δ, σ, ζ), but also on the incidence of the fundamental shock on the state variables in each country (as captured by $\boldsymbol{u}_h = \boldsymbol{R}\widetilde{\boldsymbol{f}}_{(h)}$). From Proposition 5, any empirical shock ($\widetilde{\boldsymbol{f}}$) can be expressed as a linear combination of the eigen-shocks. Therefore, the speed of convergence also varies across these empirical shocks with their incidence on the state variables in each country, reflecting the extent to which they load on eigen-shocks with slow versus fast convergence.

4 Empirical Results

We now introduce our data sources and present our main empirical results.

4.1 Data Sources

Our main dataset includes data on employment, wages, value added and bilateral mobility for 48 occupations in the United States for the period from 1980-2020. We combine data from the Current Population Survey (CPS), the Decennial Census, the American Community Survey (ACS), and the Bureau of Economic Analysis (BEA).[4]

We define occupations based on the minor occupation groups in the 1980 Standard Occupational Classification. For consistency, we combined four occupation groups that had a small number of observations in the 1980 Current Population Survey (CPS) with larger occupation groups. In the end, we have 48 occupations, plus one additional group that corresponds to unemployment. We construct data on employment and wages for each occupation for each Census year from 1980 onwards. We use the Decennial Census data for 1980, 1990 and 2000, as well as the 2010 and 2018 1-year American Community Survey data.

We compute an occupation-to-occupation mobility matrix for each Census year from 1980 onwards. We use the CPS Annual Social and Economic Supplement (ASEC), which includes questions on current and past occupations of employment. To enlarge the sample size in the CPS data, we pool data from five years surrounding each Census year.

[3] In general, these eigenvectors and eigenvalues can be complex-valued. If the initial impact is the real part of a complex eigenvector $\boldsymbol{u}_h$ ($\boldsymbol{R}\widetilde{\boldsymbol{f}} = \text{Re}(\boldsymbol{u}_h)$), then $\ln \boldsymbol{x}_{t+1} - \ln \boldsymbol{x}_t = \text{Re}(\lambda_h^t \boldsymbol{u}_h) \neq \text{Re}(\lambda_h) \cdot \text{Re}\left(\lambda_h^{t-1}\boldsymbol{u}_h\right)$. That is, the impact no longer decays at a constant rate λ_h. Instead, the complex eigenvalues introduce oscillatory motion as the dynamical system converges to the new steady state. In our empirical application, the imaginary components of $\boldsymbol{P}$'s eigenvalues are small, implying that oscillatory effects are small relative to the effects that decay exponentially.

[4] For further discussion of the data sources and definitions, see Online Appendix G.

We construct labor shares and capital stocks for each occupation using information on these variables at the industry level. We measure labor share, value added and capital stocks at the industry level using data from the Bureau of Economic Analysis (BEA) for 3-digit North American Industry Classification (NAICS) industries.

We use the cross walk from Schaller and DeCelles (2021) to concord the BEA data to each of our industries. We compute the labor share and capital stock for each occupation by taking the weighted average of the labor shares across all industries with positive employment for an occupation, using shares of the wage bill as weights.

4.2 Model Calibration

To focus on the novel implications from combining labor and capital dynamics, we assume standard values for the model's parameters from existing empirical research.

We set the elasticity of substitution between labor and capital within each occupation (ζ) equal to 1.28 based on mean estimate across occupations in Caunedo et al. (2023). This calibrated elasticity of substitution is larger than that typically assumed in the macro literature (e.g., Oberfield and Raval 2021). But our elasticity of substitution between labor and capital is within occupations, whereas the macro elasticities use variation across and within occupations. It is plausible that the elasticity of substitution is larger within than across occupations and the estimates from Caunedo et al. (2023) are estimated within occupations.

We set the elasticity of substitution across occupations (σ) equal to 1.15, which implies that capital and labor are more substitutable within occupations than across occupations ($\zeta > \sigma$). This calibrated elasticity is somewhat lower than the estimated elasticity of substitution between skilled and unskilled labor of 1.41 in Katz and Murphy (1992). But that parameter estimate uses variation within and across occupations. It is plausible that some occupations could be more suited to skilled workers than unskilled workers, and hence that the elasticity of substitution across occupations could be lower than that within occupations.

We set the migration elasticity across occupations (ρ) equal to 1.5, which is comparable to the existing estimates of 1.26-1.8 using U.S. data in Burstein et al. (2019) and 1.67 using Danish data in Traiberman (2019). We set the discount factor (β) equal to 0.95, as a central value in existing research. We set the elasticity of next period's capital stock with respect to investment (δ) equal to 0.05, which implies a steady-state ratio of investment to the capital stock that is consistent with conventional assumptions of a depreciation rate of around 5 percent.

4.3 Steady-State Gaps

We begin by using our dynamic exact-hat algebra results for the non-linear model from Proposition 1 to solve for the unobserved steady-state levels of the endogenous variables.

In Figure 1, we show the log steady state gaps ($\widetilde{x}_{ot} = \ln x_{ot} - \ln x_o^*$) for employment (panel (a)) and wages (panel (b)). We display these steady-state gaps at the beginning of our sample period (1981) and just before the COVID-19 pandemic towards the end of our sample period (2018). Each circle shows an occupation, where the size of the circles is proportional to employment in each occupation at the beginning of the sample period. The black dotted line shows the 45 degree line. The red dashed line shows the regression fit between the variables.

In each year, we find substantial gaps of both employment and wages from steady-state, consistent with the idea that adjustment towards steady-state does not occur instantaneously. We find that the slope of the regression line is less one (less than the 45-degree), which implies imperfect persistence in steady-state gaps over time, consistent with a convergence of both employment and wages towards their steady-state levels over time.

Figure 1: Log Deviations of Occupations from Steady-State in 1981 and 2018



Note: Log deviations of the actual values of variables from their steady-state values; vertical axis shows results for 2018 (before the COVID-19 pandemic); horizontal axis shows results for 1981; Panel (a) shows results for employment; Panel (b) shows results for wages; steady-state values in each year recovered by starting from the observed data in that year and using dynamic exact-hat counterfactuals in the non-linear model to solve for the implied steady-state equilibrium in the absence of any further changes in the model fundamentals (e.g., automation (α_{ot}), productivity (z_{ot}), non-wage compensating differentials (b_{ot}), investment efficiency (η_{ot}), and mobility frictions (κ_{iot}).

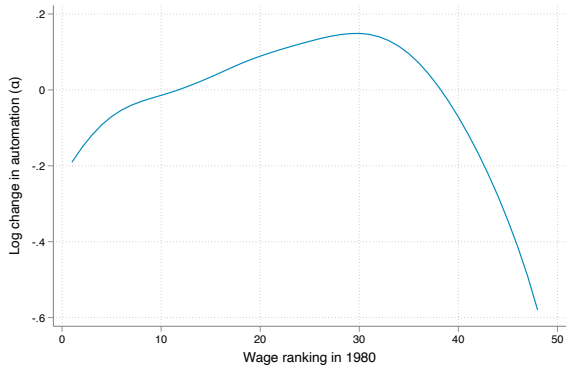
4.4 Model Inversion

We next use our model inversion results for the non-linear model from Proposition 2 to recover the fundamental shocks that rationalize the observed changes in the data.

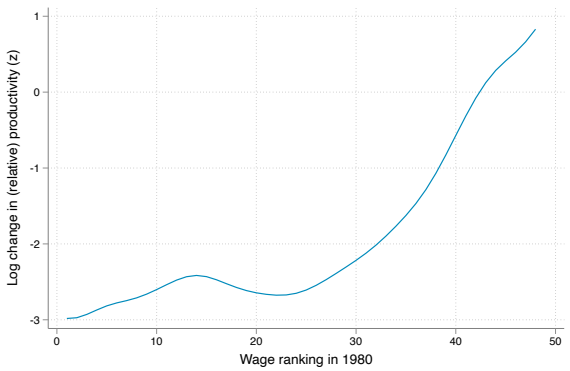
In Figure 2, we display log changes in fundamentals for each occupation from 1981-2018. We omit the shocks to our weighted measures of inward and outward mobility costs ($\Delta \ln \widetilde{\boldsymbol{\kappa}}^{in}$, $\Delta \ln \widetilde{\boldsymbol{\kappa}}^{out}$), because we find below they do not play a quantitatively relevant role in explaining the observed changes in employment shares and wages in the data. We graph the fitted values from locally-weighted linear least squares regressions of the log changes in the fundamental shocks on the initial wage ranking of the occupation in 1980 before the beginning of our sample. We allow the scale of the vertical axis to vary across each of the panels of the figure.

Figure 2: Log Changes in Fundamentals from 1981-2018 and Initial Wage Ranking of Occupations in 1980

(a) Automation (α_{ot})



(b) Productivity (z_{ot})



(c) Non-wage (b_{ot})



(d) Investment Efficiency (η_{ot})



Note: Vertical axis shows log changes in fundamentals from 1981-2018 from inverting the equilibrium conditions of the model; Horizontal axis sorts occupations by their wage ranking in 1980 (from low to high); figure shows the fitted values from locally-weighted linear least squares regressions of the log changes in fundamentals on the initial occupation wage ranking in 1980; panel (a) shows automation shocks (α_{ot}): panel (b) shows productivity shocks (z_{ot}); panel (c) shows non-wage compensating differentials shocks (b_{ot}); and panel (d) shows investment efficiency shocks (η_{ot}).

Panel (a) shows the automation shock ($\Delta \ln \alpha_{ot}$); panel (b) presents the non-neutral technical change shock ($\Delta \ln z_{ot}$); panel (c) illustrates the non-wage compensating differential shock ($\Delta \ln b_{ot}$); and panel (d) displays the investment efficiency shock ($\Delta \ln \eta_{ot}$). Consistent with existing reduced-form empirical evidence, we find that intermediate-wage occupations have experienced the greatest automation (top-left panel). In line with the literature on skill-biased technological change, we find that high-wage occupations have experienced the greatest productivity growth (top-right panel). Therefore workers in high-wage occupations have been sheltered from automation and have benefits from technological change that has raised their relative productivities,

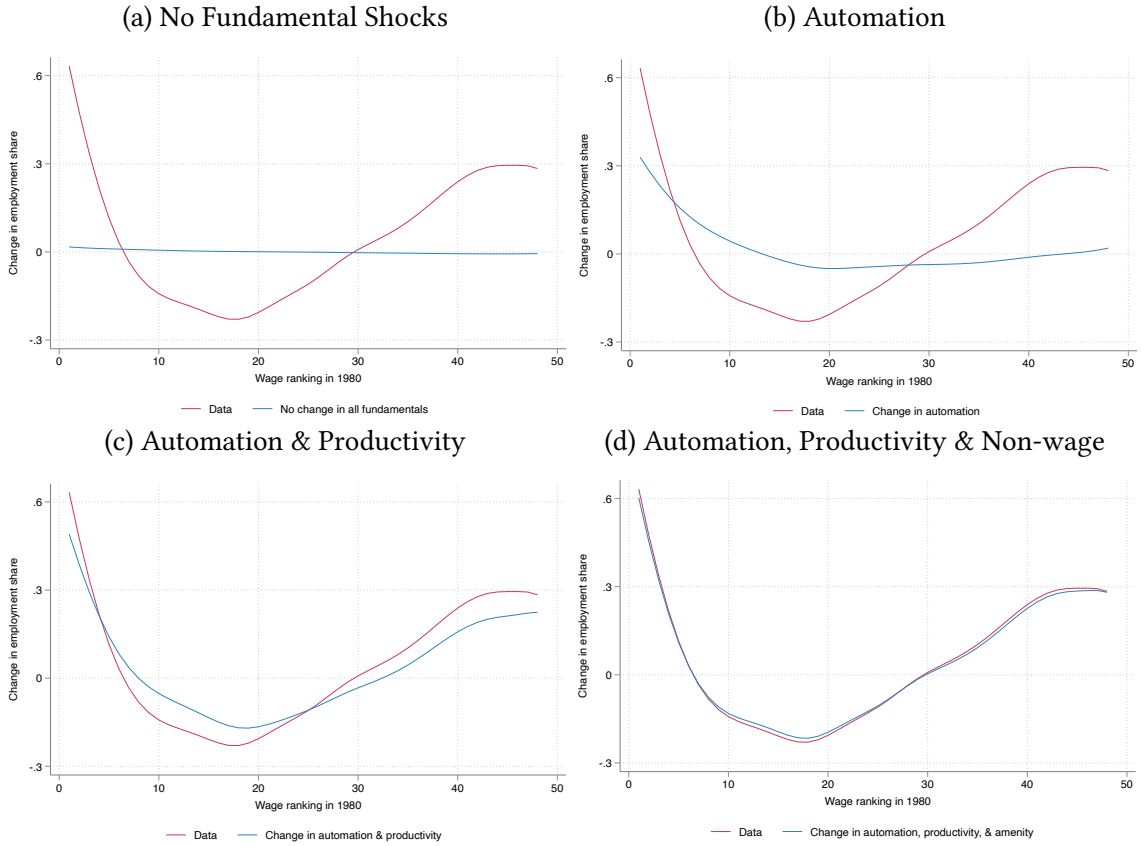
We find some evidence that intermediate occupations have become relatively less attractive in terms of non-wage compensating differentials, which could reflect for example reductions in pensions and other benefits, driven in part by a decline in unionization in intermediate-wage occupations. We also find some evidence of a decline in investment efficiency in intermediate-wage occupations, which could reflect the fact that as a larger and large share of tasks are automated in these occupations and performed by capital, there is a decline in the productivity of investment for the marginal automated task.

4.5 Counterfactuals

We next use of our measures of fundamental shocks from our model inversion in the previous subsection to undertake counterfactuals to examine the contribution of each fundamental shock to observed changes in employment and wages in the data. The counterfactuals involve solving the full general equilibrium of the model, such that a fundamental shock does not only directly affect the occupation in which it occurs, but also has indirect effects on other occupations through endogenous changes in prices and wages.

In Figure 3, we display changes in occupations shares of total employment from 1981-2018 against the initial wage ranking occupations in 1980 (where a high rank corresponds to a high wage). The red line shows the fitted values from a kernel-weighted local polynomial regression of the observed changes in employment shares on the initial wage rankings. These observed changes in employment share equal the counterfactual predictions of our model when we allow for convergence given initial conditions and all fundamental shocks. This red line for the observed data is the same across all panels of the figure. We replicate existing findings of employment polarization, with the largest increases in employment shares for occupations at the bottom and top of the initial wage distribution.

Figure 3: Changes in Employment Shares from 1981-2018 in the Data and in Model Counterfactuals Successively Adding Additional Shocks to Fundamentals



Note: Vertical axis shows change in occupation employment share (percentage point) from 1981-2018; Horizontal axis sorts occupations by their wage ranking in 1980 (from low to high); figure shows the fitted values from locally-weighted linear least squares regressions of the changes in occupation employment shares from 1981-2018 on the initial occupation wage ranking in 1980; red line is the same for all panels and shows the observed data, which is equal to the model's counterfactual predictions based on convergence to steady-state and shocks to all fundamentals; panel (a): blue line shows the model's counterfactual predictions based on convergence to steady-state; panel (b): blue line shows the model's counterfactual predictions based on convergence to steady-state and shocks to automation (α_{ot}); panel (c): blue line shows the model's counterfactual predictions based on convergence to steady-state and shocks to automation (α_{ot}) and productivity (z_{ot}); panel (d): blue line shows the model's counterfactual predictions based on convergence to steady-state and shocks to automation (α_{ot}), productivity (z_{ot}), and non-wage compensating differentials (b_{ot}).

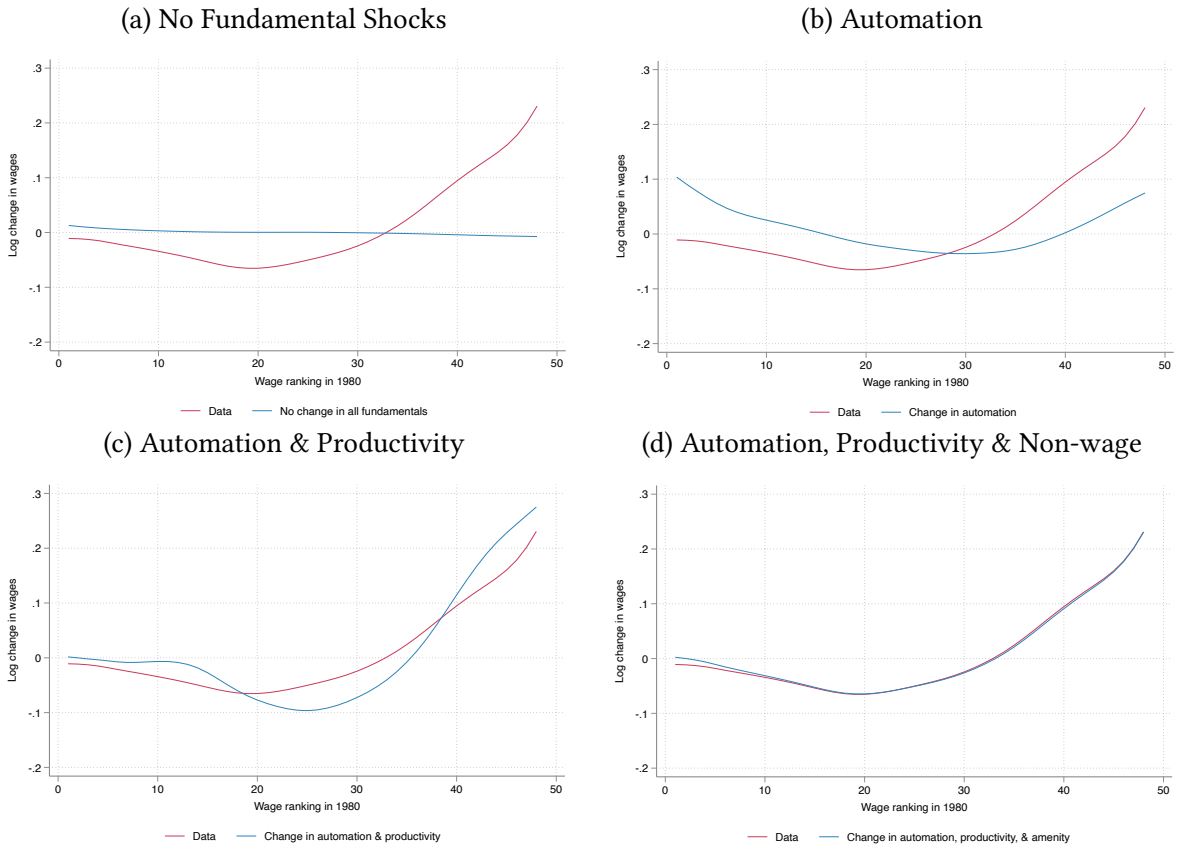
The blue line shows the fitted values from a kernel-weighted local polynomial regression of counterfactual changes in employment shares on the initial wage rankings, where these counterfactuals progressively add additional determinants of changes in employment shares over time. Panel (a) includes only convergence given initial conditions; Panel (b) incorporates convergence given initial conditions and automation shocks; Panel (c) admits convergence given initial conditions, automation shocks and non-neutral technical change; and Panel (d) encompasses convergence given initial conditions, automation shocks, non-neutral technical change, and non-wage compensating differentials.

We find that convergence given initial conditions makes a relatively small contribution towards explaining the observed changes in employment shares, as reflected in the relatively flat blue line in Panel (a). We find that the automation shock makes a substantial contribution towards explaining the polarization of employment in low-wage occupations, as reflected in the emergence of a downward-slope in the blue line in Panel (b). When we combine both automation and non-neutral technical change, the model successfully captures the polarization of employment in both low and high-wage occupations, as reflected in the clear U-shaped pattern in the blue line in panel (c). A key implication is that it is the combination of automation and non-neutral technical change that drives the observed polarization of employment at the top and bottom of the wage distribution.

Once we simultaneously incorporate automation, non-neutral technical change and non-wage compensating differentials in panel (d), we find that the counterfactual predictions of our model align closely with the observed data, as reflected in the blue line lying close to the red line throughout the entire of the wage distribution. Therefore, we find that the omitted fundamental shocks of investment efficiency and our weighted measures of inward and outward mobility frictions make a relatively small contribution towards the observed employment polarization. In principle, these other forces could have played an important role in explaining changes in observed employment shares, but we find that they are dominated by the contributions of automation and non-neutral technical change.

In Figure 4, we display log changes in wages from 1981-2018 against the initial wage ranking occupations in 1980 (where a high rank corresponds to a high wage). The structure is the same as in the previous figure. The red line is the same in all panels and shows the fitted values from a kernel-weighted local polynomial regression of observed wage growth on initial wage rankings. These observed changes in log wages equal the counterfactual predictions of our model when we allow for convergence given initial conditions and all fundamental shocks. This red line for the observed data is the same across all panels of the figure. We replicate existing findings for wage inequality, with some evidence of wage polarization at both ends of the wage distribution, but the greatest wage growth at the top of the wage distribution.

Figure 4: Log Changes in Wages from 1981-2018 in the Data and in Model Counterfactuals Adding Shocks to Fundamentals



Note: Vertical axis shows log changes in wages from 1981-2018; Horizontal axis sorts occupations by their wage ranking in 1980 (from low to high); figure shows the fitted values from locally-weighted linear least squares regressions of the log changes in wages from 1981-2018 on the initial occupation wage ranking in 1980; red line is the same for all panels and shows the observed data, which is equal to the model's counterfactual predictions based on convergence to steady-state and shocks to all fundamentals; panel (a): blue line shows the model's counterfactual predictions based on convergence to steady-state; panel (b): blue line shows the model's counterfactual predictions based on convergence to steady-state and shocks to automation (α_{ot}); panel (c): blue line shows the model's counterfactual predictions based on convergence to steady-state and shocks to automation (α_{ot}) and productivity (z_{ot}); panel (d): blue line shows the model's counterfactual predictions based on convergence to steady-state and shocks to automation (α_{ot}), productivity (z_{ot}), and non-wage compensating differentials (b_{ot}).

We again find that convergence given initial conditions makes a relatively small contribution towards explaining the observed changes in the data, as reflected in the relatively flat blue line in Panel (a). We find that adding the automation shock generates wage polarization as the top and bottom of the wage distribution, as reflected in the emergence of a U-shaped blue line in Panel (b). When we combine both automation and non-neutral technical change, we generate the pattern of stronger wage growth at the top than at the bottom of the wage distribution, as reflected in the blue line exhibiting a stronger upward slope at the top of the wage distribution in Panel (c). Therefore, we again find that it is the combination of automation and non-neutral technical change that is needed to match the patterns in the data, mirroring our findings for

employment shares above.

Once we simultaneously incorporate automation, non-neutral technical change and non-wage compensating differentials in panel (d), our model's predictions closely align with the observed wage changes in the data, as reflected in the blue line lying virtually on top of the red line throughout the wage distribution. Therefore, we find that the omitted fundamental shocks of investment efficiency and our weighted measures of inward and outward mobility frictions are again relatively unimportant for explaining the observed log changes in wages. In another counterfactual world, these other forces could could have played a leading role in explaining changes in wage inequality, but we instead find that they are dominated in reality by the contributions from automation and non-neutral technical change.

4.6 Efficiency Versus Equity Trade-off

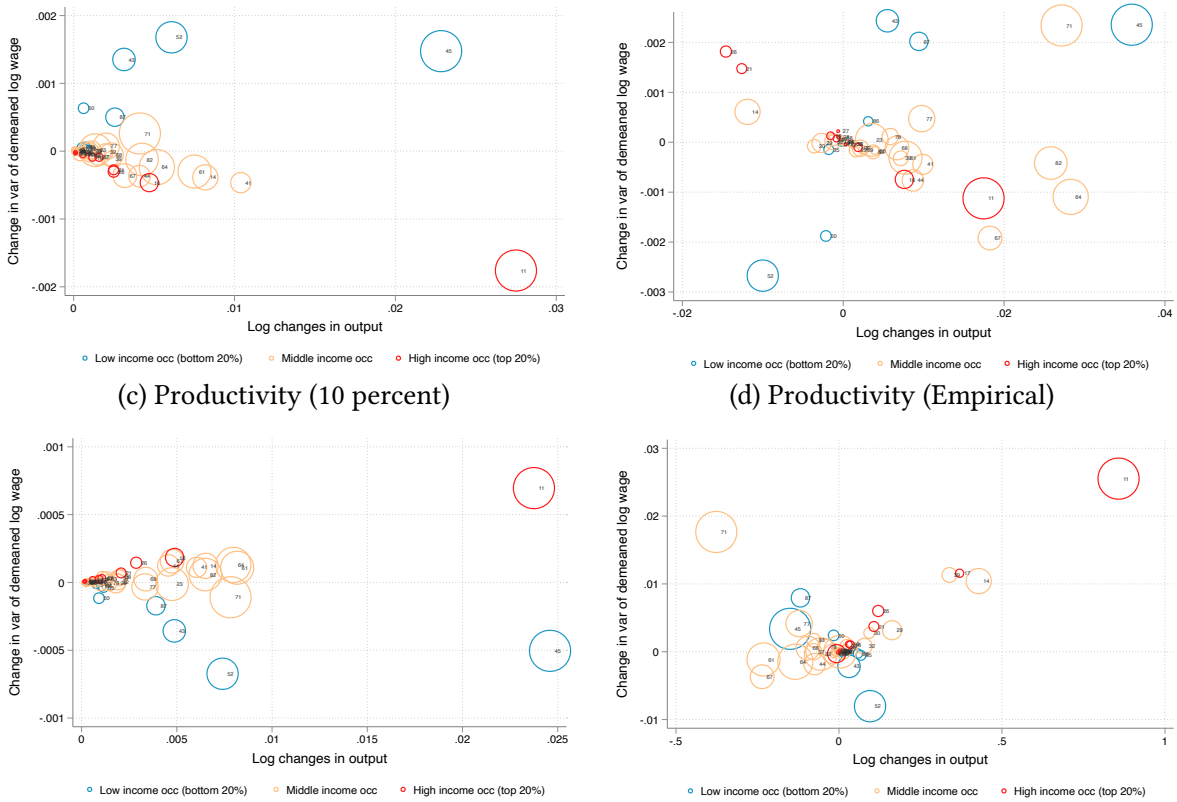
Our model provides a framework for evaluating the potential tradeoff between the efficiency and equity effects of automation and non-neutral technological change. Innovation shifts the production possibility frontier outwards the aggregate size of the pie. But to the extent that these technological advances occur unevenly across occupations they typically have income distributional consequences for workers, creating winners and losers, depending on the extent to which they substitute for workers in each occupation and reduce consumer prices.

Our quantitative model provides a suitable framework for assessing this efficiency-equity trade-off in an empirically realistic setting with many occupations, dynamics from worker mobility capital accumulation, and a rich specification of technology that encompasses both automation and non-neutral technical change. In Figure 5, we display the results of counterfactuals shocking either automation (α_{ot}) or productivity (z_{ot}). Panel (a) shows results for a 10 percent automation shock; Panel (b) shows results for the empirical automation shock; Panel (c) shows results for a 10 percent productivity shock; Panel (d) shows results for the empirical productivity shock. In each case, we shock each occupation individually one at a time, and solve for general equilibrium. The vertical axis displays the change in the variance of log wages as a measure of equity. The horizontal axis denotes the change in aggregate output as a measure of efficiency, where recall that the aggregate price index is our numeraire. Each occupation is indicated by a circle and its associated 2-digit code, where the size of the circles is proportional to the initial employment size of the occupations in 1980.

Figure 5: Counterfactuals for Output-Inequality Trade-off

(a) Automation (10 percent)

(b) Automation (Empirical)



(c) Productivity (10 percent)

(d) Productivity (Empirical)

Note: Counterfactual predictions of the model from a shock to fundamentals; vertical axis in each panel shows counterfactual change in the variance of demeaned log wages; horizontal axis shows log changes in aggregate output; blue circles show low-wage occupations (bottom 20 percent of 1980 wage ranking); orange circles show middle-income occupations (middle 60 percent of 1980 wage ranking); red circles show high-wage occupations (top 20 percent of 1980 wage ranking); size of circles is proportional to 1980 employment of occupation; Panel (a) shows the results from a 10 percent automation shock in each occupation separately; Panel (b) shows the results from the empirical automation shock from 1981-2018 in each occupation separately; panel (c) shows the results from a 10 percent productivity shock in each occupation separately; panel (d) shows the results the empirical productivity shock in each occupation separately.

In panel (a), we find an intuitive pattern, in which automation in low-wage occupations (bottom 20 percent of the initial wage distribution) increases wage inequality, whereas automation in high-wage occupations (top 20 percent of the initial wage distribution) reduces wage inequality. However, these is considerable heterogeneity within each of these groups in the increase in aggregate output achieved by a given amount of automation, in part related to the initial employment size of the occupations. Automation in some occupations brings large gains in aggregate output with little increase or a reduction in wage inequality (e.g., Officials and Administrators (11)), while automation in other occupations has only a modest impact on aggregate output at the expense of a substantial increase in wage inequality (e.g., Service Occupations (52)). Comparing panels (a) and (b), we find substantial changes in the impact of automation on both wage inequality and

aggregate output once we take into account the empirical size of the automation shock, because the empirical size of this shock varies substantially across occupations, with intermediate-wage occupations experiencing the greatest automation.

In panel (c), we also an intuitive pattern for non-neutral technical change, although one that differs substantially from that for automation. Productivity growth in low-wage occupations (bottom 20 percent of the initial wage distribution) reduces wage inequality, whereas productivity growth in high-wage occupations (top 20 percent of the initial wage distribution) increases wage inequality. Again these is considerable heterogeneity within each of these groups in the increase in aggregate output achieved by a given amount of productivity growth, in part related to the initial employment size of the occupations. Productivity growth in some occupations brings large gains in aggregate output with little increase or a reduction in wage inequality (e.g., Supervisors and Administrative Support (45)), while productivity growth in other occupations has only a modest impact on aggregate output at the expense of a substantial increase in wage inequality (e.g., Physicians and Dentists (26)). Comparing panels (c) and (d), we find substantial changes in the impact of productivity growth on both wage inequality and aggregate output once we take into account the empirical magnitude of productivity growth, because the empirical magnitude of this productivity growth varies substantially across occupations, with high-wage occupations experiencing the greatest productivity growth.

To the extent that policy makers are concerned about equity, these empirical findings provide guidance as to how policies intended to influence the extent of automation or non-neutral technical change can achieve equity objectives at relatively low cost in terms of aggregate output.

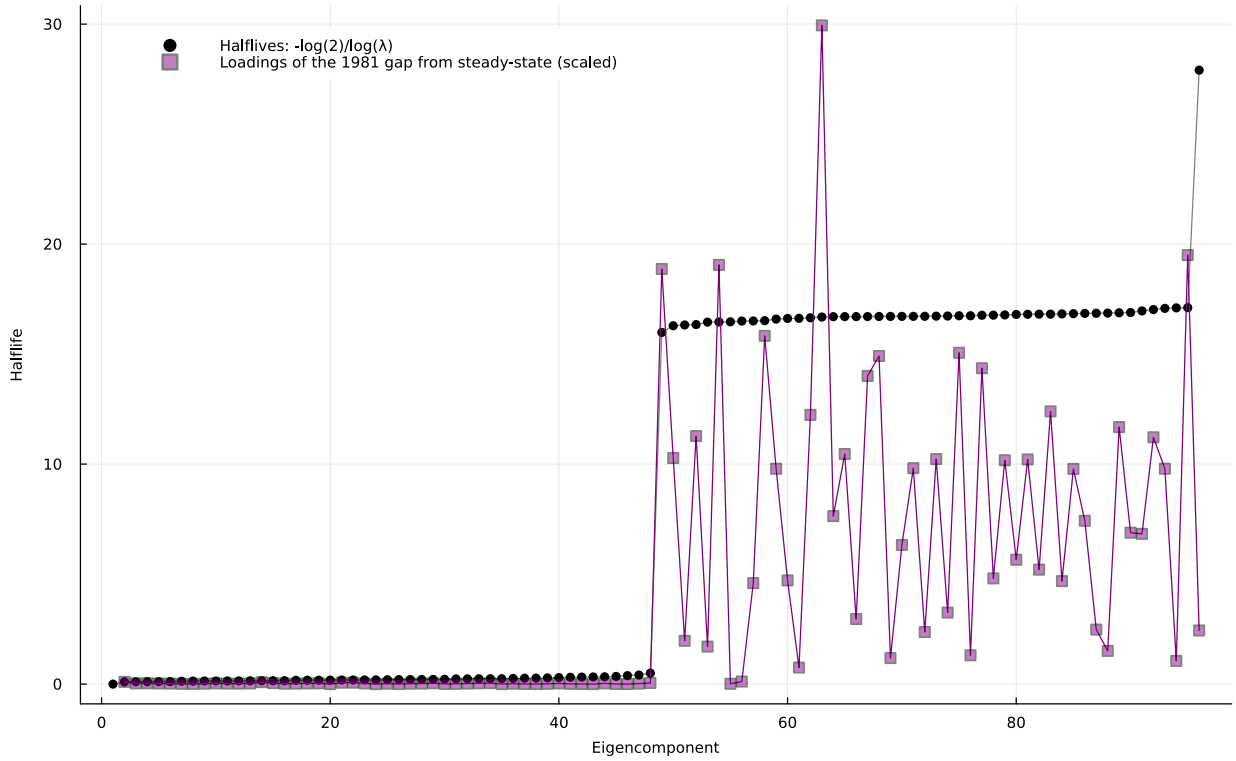
4.7 Speed of Convergence

One of our findings in the previous sections is that convergence given initial conditions makes a relatively small contribution to explaining the observed employment and wage polarization in the data. We how show that one of the reasons for this finding is that the rate of convergence to steady-state is relatively high, because of the high mobility of workers between occupations in the data.

Using Propositions 4-6, we compute half-lives of convergence to steady state as determined by the eigenvalues of the transition matrix. In Figure 6, we show these half-lives (solid black line with circle markers) for the entire spectrum of 2O eigencomponents, sorted by increasing half-life. Each eigencomponent corresponds to an eigenshock for which the initial impact of the shock on the state variables is equal to an eigenvector of the transition matrix ($\boldsymbol{u_h} = \boldsymbol{R}\widetilde{\boldsymbol{f}}_{(h)}$). We display results based on the transition matrix ($\boldsymbol{P}$) computed using the implied steady-state worker mobility matrices at the beginning of our sample period in 1981. We compute these implied steady-state matrices using using our dynamic exact-hat algebra results from Proposition

1.

Figure 6: Half-Lives of Convergence Towards Steady-State



Note: Black line shows half lives of convergence to steady state (left vertical axis) based on the transition matrix for 1981; purple line shows the loadings of the log employment gap from steady-state in 1981 on the eigencomponents (right vertical axis); left vertical axis shows the half life in years (the time taken to close half of the gap to steady-state); right vertical axis shows loadings on the eigencomponents; horizontal axis sorts the eigencomponents in terms of increasing half lives.

Three features are particularly noteworthy. First, the speed of convergence to steady-state is typically relatively fast, with a half-life across the entire spectrum of eigenshocks of typically less than 20 years. This finding contrasts with evidence from dynamic discrete choice models of worker mobility across locations, and reflects the fact that in the observed data worker mobility across occupations is substantially larger than worker mobility across locations. Second, there is substantial heterogeneity in the speed of convergence across eigenshocks, with the half-life of convergence varying from less than one year to just over 20 years. Third, the correlation between the gaps of the labor and capital state variables from steady-state plays a key role in determining the speed of convergence to steady-state.

The first O eigencomponents in Figure 6 with low half-lives reflect shocks to the state variables in which the gaps of labor and capital from steady-state are negatively correlated with one another. For these shocks, when capital is above steady-state, labor tends to be below steady-state. Since capital and labor are substitutes within occupations (an elasticity of substitution greater than one) and worker mobility across occupations is high, the economy can adjust rela-

tively rapidly to these shocks. The second O eigencomponents in Figure 6 with high half lives reflect shocks to the state variables in which the gaps of labor and capital from steady-state are positively correlated with one another. For these shocks, when capital is above steady-state, labor tends to be above steady-state too. Since capital and labor are substitutes within occupations, adjustment to these shocks requires both worker mobility and adjustment costs for investment to be incurred, which slows down the speed of convergence (a larger half life) towards steady-state.

Therefore, the speed of the economy's adjustment to empirical shocks to fundamentals can vary substantially, depending on the extent to which these empirical shocks load on eigencomponents with either low or high half-lives. The purple line in Figure 6 shows the loadings of the log employment gap from steady-state in 1981 on the eigencomponents. We find than these loadings are more concentrated on the second O eigencomponents with high half-lives than the first O eigencomponents with low half-lives, implying a half-life taking a number of years for these initial gaps of occupation employment from steady-state.

5 Conclusions

A key research and public policy question is the effects of new technologies, such as automation and artificial intelligence, on income distribution and welfare. These effects on income distribution and welfare are likely to emerge gradually over time, because it takes time for workers to reallocate in response to labor demand shocks, and it takes time for complementary investments in capital to occur to realize the potential of these new technologies.

We develop a quantitative theoretical model with dynamic labor mobility and capital accumulation that remains tractable even with a high-dimensional state space. At each point in time, workers supply labor in their existing occupation. They also make forward-looking decisions whether to incur mobility costs to reallocate to other occupations, subject to idiosyncratic shocks to these mobility costs. Firms produce output by performing a set of procedures that are associated with occupations, where each procedure involves tasks that can be performed either by labor or capital. Capital and labor are perfect substitutes at the task level, after adjusting for idiosyncratic shocks to the productivity of capital and labor in performing individual tasks. Aggregating across tasks, we obtain a production technology for each procedure that takes the constant elasticity of substitution (CES) form. Capital owners make forward-looking decisions about investment in each procedure subject to adjustment costs.

We implement our quantitative model using U.S. data on 48 occupations for the period from 1981-2018, before the beginning of the COVID-19 pandemic. Given observed data on wages, employment, capital payments, capital stocks and worker mobility for an economy that is somewhere along the transition path to a steady-state, we invert the equilibrium conditions of the

model and recover unique values of the fundamental shocks that exactly rationalize the observed data as an equilibrium of the model. These fundamental shocks include automation, non-neutral technical change, non-wage compensating differentials, investment efficiency and bilateral mobility costs. We use of our measures of fundamental shocks from this model inversion to undertake counterfactuals to examine the contribution of convergence given initial conditions and each of these fundamental shocks to observed changes in employment and wages in the data.

Consistent with existing reduced-form empirical evidence, we find that intermediate-wage occupations have experienced the greatest automation. In line with the literature on skill-biased technological change, we find that high-wage occupations have experienced the greatest productivity growth. Implementing our model counterfactuals, we find that neither automation alone, nor non-neutral technical change alone, is able to fully explain the observed patterns in the data, which include employment polarization for occupations at the top and bottom of the wage distribution, and a the greater increase in wages for occupations at the top of the wage distribution. Instead, we show that it is the combination of automation and non-neutral technical change that is needed to fully match these patterns in the data. In comparison, we find that non-wage compensating differentials, investment efficiency and mobility costs make relatively minor contributions to observed changes in employment and wages.

We use our model counterfactuals to assess the extent to which gains in aggregate output from technological change come at the expense of increases in wage inequality. Automation in low-wage occupations (bottom 20 percent of the initial wage distribution) increases wage inequality, whereas automation in high-wage occupations (top 20 percent of the initial wage distribution) reduces wage inequality. In contrast, productivity growth reduces wage inequality when it occurs in low-wage occupations and increases wage inequality when it occurs in high-wage occupations. We find considerable heterogeneity across occupations in the extent to which a given amount of automation or productivity growth affects wage inequality and aggregate output, in part related to the initial employment size of occupations. To the extent that policy makers are concerned about equity, these empirical findings provide guidance as to how policies intended to influence the direction of automation or non-neutral technical change can achieve equity objectives at relatively low cost in terms of aggregate output.

We find that convergence given initial conditions makes a relatively small contribution to explaining the observed employment and wage polarization in the data. We how show that one of the reasons for this finding is that the rate of convergence to steady-state is relatively high, because of the high observed mobility of workers between occupations in the data. Nonetheless, the empirical gaps of occupation employment from steady-state at the beginning of our sample period imply half-lifes of convergence to steady-state of a number of years.

References

Acecmoglu, D. and D. Autor (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, ed. by D. Card and O. Ashenfelter, Amsterdam: Elsevier: North-Holland, vol. 4, chap. 12, 1043–1171.

Acemoglu, D. (2002): "Directed Technical Change," *Review of Economic Studies*, 69, 781–809.

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo (2022): "Artificial Intelligence and Jobs - Evidence from Online Vacancies," *Journal of Labor Economics*, S293–S340.

Acemoglu, D., G. Gancia, and G. Zilibotti (2015): "Offshoring and Directed Technical Change," *American Economic Journal: Macro*, 7, 84–112.

Acemoglu, D. and S. Johnson (2023): *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*, New York: Public Affairs.

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

——— (2020): "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 128, 2188–2244.

——— (2022): "Tasks, Automation and the Rise in US Wage Inequality," *Econometrica*, 90, 1973–2016.

Artuç, E., S. Chaudhuri, and J. McLaren (2010): "Trade Shocks and Labor Adjustment: A Structural Empirical Approach," *American Economic Review*, 100, 1008–45.

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

Autor, D. H. and D. Dorn (2013): "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 103, 1553–97.

Autor, D. H., L. F. Katz, and M. S. Kearney (2008): "Trends in U.S. Wage Inequality: Revising the Revisionists," *Review of Economics and Statistics*, 90, 300–323.

Babina, T., A. Fedyk, A. He, and J. Hodson (2022): "Artificial Intelligence, Firm Growth, and Product Innovation," University of California Berkeley.

Burstein, A., E. Morales, and J. Vogel (2019): "Changes in Between-Group Inequality: Computers, Occupations, and International Trade," *American Economic Journal: Macroeconomics*, 11, 348–400.

Caliendo, L., M. Dvorkin, and F. Parro (2019): "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock," *Econometrica*, 87, 741–835.

Caunedo, J., D. Jaume, and E. Keller (2023): "Occupational Exposure to Capital-embodied Technical Change," *American Economic Review*, 113, 1642–85.

Costinot, A. and I. Werning (2023): "Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation," *Review of Economic Studies*, 90, 2261–2291.

Dix-Carneiro, R. (2014): "Trade Liberalization and Labor Market Dynamics," *Econometrica*, 82, 825–885.

Goldin, C. and L. F. Katz (2010): *The Race between Education and Technology*, Cambridge MA: Harvard University Press.

Hanlon, W. (2015): "Necessity is the Mother of Invention: Input Supplies and Directed Technical Change," *Econometrica*, 83, 67–100.

Hémous, D. and M. Olsen (2021): "Directed Technical Change in Labor and Environmental Economics," *Annual Review of Economics*, 13, 571–597.

Hercowitz, Z. and M. Sampson (1991): "Output Growth, the Real Wage, and Employment Fluctations," *American Economic Review*, 81, 1215–1237.

Humlum, A. (2021): "Robot Adoption and Labor Market Dynamics," Chicago Booth School of Business.

Kaplan, G. and G. L. Violante (2014): "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82, 1199–1239.

Katz, L. F. and K. M. Murphy (1992): "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 107, 35–78.

Kleinman, B., E. Liu, and S. J. Redding (2023): "Dynamic Spatial General Equilibrium," *Econometrica*, 91, 385–424.

Liu, E. and S. Ma (2023): "Innovation Networks and R&D Allocation," *NBER Working Paper*, princeton University, mimeo.

Loebbing, J. (2022): "An Elementary Theory of Directed Technical Change and Wage Inequality," *Review of Economic Studies*, 89, 411–451.

Lucas, R. E. and E. Prescott (1971): "Investment Under Uncertainty," *Econometrica*, 39, 659–681.

McElheran, K., J. F. Li, E. Brynjolfsson, Z. Kroff, E. Dinlersoz, L. S. Foster, and N. Zolas (2023): "AI Adoption in America: Who, What and Where," *NBER Working Paper*, 31788.

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

Moll, B. (2014): "Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation," *American Economic Review*, 104, 3186–3221.

Oberfield, E. and D. Raval (2021): "Micro Data and Macro Technology," *Econometrica*, 89, 703–732.

Schaller, Z. and P. DeCelles (2021): "Weighted Crosswalks for NAICS and SIC Industry Codes,"

Ann Arbor, MI: Inter-university Consortium for Political and Social Research.

Traiberman, S. (2019): "Occupations and Import Competition: Evidence from Denmark," *American Economic Review*, 109, 4260–4301.

Uhlig, H. (1999): "A Toolkit for Analyzing Non-linear Dynamic Stochastic Models Easily," in *Computational Methods for the Study of Dynamic Economies*, ed. by R. Marimon and A. Scott, New York: Oxford University Press, 30–61.